

EXHIBIT 25

**IN THE UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF TEXAS
WACO DIVISION**

Sonos, Inc., §
§
v. Plaintiff, § No. 6:20-cv-881-ADA
§
Google LLC, §
§
Defendant. §

DECLARATION OF KEVIN C. ALMEROOTH

I, Kevin C. Almeroth, hereby declare as follows:

I. INTRODUCTION

1. This Declaration is based upon my own personal knowledge.
2. If I am called upon to testify regarding this Declaration, I would testify competently and truthfully to the matters discussed herein.
3. I understand that in the above-captioned action between Plaintiff Sonos, Inc. (“Sonos”) and Defendant Google LLC, Sonos has alleged that Google infringes certain claims of U.S. Patent Nos. 9,344,206 (the “‘206 patent”), 10,469,966 (the “‘966 patent”), and 10,848,885 (the “‘885 patent”) (collectively referred to herein as the “Zone Scene Patents”).
4. I understand that the Zone Scene Patents are all part of the same patent family that originated with U.S. Provisional Application No. 60/825,407 (“the ‘407 Provisional”), which was filed on September 12, 2006.
5. It is my understanding that the asserted claims for each of the Zone Scene Patents are as follows:
 - ‘206 patent – claims 1-5, 7, 10-19

- ‘966 patent – claims 1-4, 6-12, 14-20
- ‘885 patent – claims 1-3, 5-10, 12-17, and 19-20

6. The term “data network” appears in dependent claims 7 and 15 of the ‘966 patent and all asserted claims of the ‘885 patent. I have been asked by Sonos to provide my opinions regarding the construction of this claim term.

7. In forming my opinions, I have reviewed the Zone Scene Patents and their respective file histories, as well as the ‘407 Provisional to which the Zone Scene Patents claim priority. I note that, while the specifications of the Zone Scene Patents are not identical, the disclosure of each Zone Scene Patent is substantially the same.

8. I reserve the right to supplement or clarify the opinions set forth herein, and if I am requested to do so, to provide additional opinions regarding the asserted claims of the Zone Scene Patents.

II. MY BACKGROUND

9. I am currently a Professor Emeritus in the Department of Computer Science at the University of California, Santa Barbara (UCSB). While active at UCSB, I held faculty appointments and was a founding member of the Computer Engineering (CE) Program, Media Arts and Technology (MAT) Program, and the Technology Management Program (TMP). I also served as the Associate Director of the Center for Information Technology and Society (CITS) from 1999 to 2012. I have been a faculty member at UCSB since July 1997.

10. I hold three degrees from the Georgia Institute of Technology: (1) a Bachelor of Science degree in Information and Computer Science (with minors in Economics, Technical Communication, and American Literature) earned in June 1992; (2) a Master of Science degree in Computer Science (with specialization in Networking and Systems) earned in June 1994; and (3)

a Doctor of Philosophy (Ph.D.) degree in Computer Science (Dissertation Title: Networking and System Support for the Efficient, Scalable Delivery of Services in Interactive Multimedia System, minor in Telecommunications Public Policy) earned in June 1997. During my education, I have taken a wide variety of courses as demonstrated by my minor. My undergraduate degree also included a number of courses more typical of a degree in electrical engineering including digital logic, signal processing, and telecommunications theory.

11. One of the major concentrations of my research has been the delivery of multimedia content and data between computing devices, including various network architectures. In my research, I have studied large-scale content delivery systems, and the use of servers located in a variety of geographic locations to provide scalable delivery to hundreds or thousands of users simultaneously. I have also studied smaller-scale content delivery systems in which content is exchanged between individual computers and portable devices. My work has emphasized the exchange of content more efficiently across computer networks, including the scalable delivery of content to many users, mobile computing, satellite networking, delivering content to mobile devices, and network support for data delivery in wireless networks.

12. In 1992, the initial focus of my research was on the provision of interactive functions (e.g., VCR-style functions like pause, rewind, and fast-forward) for near video-on-demand systems in cable systems; in particular, how to aggregate requests for movies at a cable head-end and then how to satisfy a multitude of requests using one audio/video stream broadcast to multiple receivers simultaneously. This research has continually evolved and resulted in the development of techniques to scalably deliver on-demand content, including audio, video, web documents, and other types of data, through the Internet and over other types of networks, including over cable systems, broadband telephone lines, and satellite links.

13. An important component of my research has been investigating the challenges of communicating multimedia content, including video, between computers and across networks including the Internet. Although the early Internet was used mostly for text-based, non-real time applications, the interest in sharing multimedia content, such as video, quickly developed. Multimedia-based applications ranged from downloading content to a device to streaming multimedia content to be instantly used. One of the challenges was that multimedia content is typically larger than text-only content, but there are also opportunities to use different delivery techniques since multimedia content is more resilient to errors. I have worked on a variety of research problems and used a number of systems that were developed to deliver multimedia content to users. One content-delivery method I have researched is the one-to-many communication facility called “multicast,” first deployed as the Multicast Backbone, a virtual overlay network supporting one-to-many communication. Multicast is one technique that can be used on the Internet to provide streaming media support for complex applications like video-on-demand, distance learning, distributed collaboration, distributed games, and large-scale wireless communication. The delivery of media through multicast often involves using Internet infrastructure, devices and protocols, including protocols for routing and TCP/IP.

14. Starting in 1997, I worked on a project to integrate the streaming media capabilities of the Internet together with the interactivity of the web. I developed a project called the Interactive Multimedia Jukebox (IMJ). Users would visit a web page and select content to view. The content would then be scheduled on one of a number of channels, including delivery to students in Georgia Tech dorms delivered via the campus cable plant. The content of each channel was delivered using multicast communication.

15. In the IMJ, the number of channels varied depending on the capabilities of the server including the available bandwidth of its connection to the Internet. If one of the channels was idle, the requesting user would be able to watch their selection immediately. If all channels were streaming previously selected content, the user's selection would be queued on the channel with the shortest wait time. In the meantime, the user would see what content was currently playing on other channels, and because of the use of multicast, would be able to join one of the existing channels and watch the content at the point it was currently being transmitted.

16. The IMJ service combined the interactivity of the web with the streaming capabilities of the Internet to create a jukebox-like service. It supported true Video-on-Demand when capacity allowed, but scaled to any number of users based on queuing requested programs. As part of the project, we obtained permission from Turner Broadcasting to transmit cartoons and other short-subject content. We also connected the IMJ into the Georgia Tech campus cable television network so that students in their dorms could use the web to request content and then view that content on one of the campus's public access channels.

17. More recently, I have also studied issues concerning how users choose content, especially when considering the price of that content. My research has examined how dynamic content pricing can be used to control system load. By raising prices when systems start to become overloaded (*i.e.*, when all available resources are fully utilized) and reducing prices when system capacity is readily available, users' capacity to pay as well as their willingness can be used as factors in stabilizing the response time of a system. This capability is particularly useful in systems where content is downloaded or streamed on-demand to users.

18. As a parallel research theme, starting in 1997, I began researching issues related to wireless devices and sensors. In particular, I was interested in showing how to provide greater

communication capability to “lightweight devices,” *i.e.*, small form-factor, resource-constrained (*e.g.*, CPU, memory, networking, and power) devices. Starting in 1998, I published several papers on my work to develop a flexible, lightweight, battery-aware network protocol stack. The lightweight protocols we envisioned were similar in nature to protocols like Bluetooth, Universal Plug and Play (UPnP) and Digital Living Network Alliance (DLNA).

19. From this initial work, I have made wireless networking—including ad hoc, mesh networks and wireless devices—one of the major themes of my research. One topic includes developing applications for mobile devices, for example, virally exchanging and tracking “coupons” through “opportunistic contact” (*i.e.*, communication with other devices coming into communication range with a user). Other topics include building network communication among a set of mobile devices unaided by any other kind of network infrastructure. Yet another theme is monitoring wireless networks, in particular different variants of IEEE 802.11 compliant networks, to (1) understand the operation of the various protocols used in real-world deployments, (2) use these measurements to characterize use of the networks and identify protocol limitations and weaknesses, and (3) propose and evaluate solutions to these problems.

20. Protecting networks, including their operation and content, has been an underlying theme of my research almost since the beginning of my research career. Starting in 2000, I have been involved in several projects that specifically address security, network protection, and firewalls. After significant background work, a team on which I was a member successfully submitted a \$4.3M grant proposal to the Army Research Office (ARO) at the Department of Defense to propose and develop a high-speed intrusion detection system. Key aspects of the system included associating streams of packets and analyzing them for viruses and other malware. Once the grant was awarded, we spent several years developing and meeting the milestones of the

project. A number of my students worked on related projects and published papers on topics ranging from intrusion detection to developing advanced techniques to be incorporated into firewalls. I have also used firewalls, including their associated malware detection features, in developing techniques for the classroom to ensure that students are not distracted by online content.

21. Recent work ties some of the various threads of my past research together. I have investigated content delivery in online social networks and proposed reputation management systems in large-scale social networks and marketplaces. On the content delivery side, I have looked at issues of caching and cache placement, especially when content being shared and the cache has geographical relevance. We were able to show that effective caching strategies can greatly improve performance and reduce deployment costs. Our work on reputation systems showed that reputations have economic value, and as such, creates a motivation to manipulate reputations. In response, we developed a variety of solutions to protect the integrity of reputations in online social networks. The techniques we developed for content delivery and reputation management were particularly relevant in peer-to-peer communication.

22. As an important component of my research program, I have been involved in the development of academic research into available technology in the market place. One aspect of this work is my involvement in the Internet Engineering Task Force (IETF). The IETF is a large and open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet. I have been involved in various IETF groups including many content delivery-related working groups like the Audio Video Transport (AVT) group, the MBone Deployment (MBONED) group, Source Specific Multicast (SSM) group, the Inter-Domain Multicast Routing (IDMR) group, the

Reliable Multicast Transport (RMT) group, the Protocol Independent Multicast (PIM) group, etc. I have also served as a member of the Multicast Directorate (MADDOGS), which oversaw the standardization of all things related to multicast in the IETF. Finally, I was the Chair of the Internet2 Multicast Working Group for seven years.

23. My involvement in the research community extends to leadership positions for several academic journals and conferences. I am the co-chair of the Steering Committee for the ACM Network and System Support for Digital Audio and Video (NOSSDAV) workshop and on the Steering Committees for the International Conference on Network Protocols (ICNP), ACM Sigcomm Workshop on Challenged Networks (CHANTS), and IEEE Global Internet (GI) Symposium. I have served or am serving on the Editorial Boards of IEEE/ACM Transactions on Networking, IEEE Transactions on Mobile Computing, IEEE Network, ACM Computers in Entertainment, AACE Journal of Interactive Learning Research (JILR), and ACM Computer Communications Review. I have co-chaired a number of conferences and workshops including the IEEE International Conference on Network Protocols (ICNP), IEEE Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON), International Conference on Communication Systems and Networks (COMSNETS), IFIP/IEEE International Conference on Management of Multimedia Networks and Services (MMNS), the International Workshop On Wireless Network Measurement (WiNMee), ACM Sigcomm Workshop on Challenged Networks (CHANTS), the Network Group Communication (NGC) workshop, and the Global Internet Symposium, and I have served on the program committees for numerous conferences.

24. Furthermore, in the courses I taught at UCSB, a significant portion of my curriculum covered aspects of the Internet and network communication including the physical and data link layers of the Open System Interconnect (OSI) protocol stack, and standardized protocols

for communicating across a variety of physical media such as cable systems, telephone lines, wireless, and high-speed Local Area Networks (LANs). The courses I have taught also cover most major topics in Internet communication, including data communication, multimedia encoding, and mobile application design. My research and courses have covered a range of physical infrastructures for delivering content over networks, including cable, Integrated Services Digital Network (ISDN), Ethernet, Asynchronous Transfer Mode (ATM), fiber, and Digital Subscriber Line (DSL). For a complete list of courses I have taught, see my curriculum vitae (CV).

25. In addition, I co-founded a technology company called Santa Barbara Labs that was working under a sub-contract from the U.S. Air Force to develop very accurate emulation systems for the military's next generation internetwork. Santa Barbara Labs' focus was in developing an emulation platform to test the performance characteristics of the network architecture in the variety of environments in which it was expected to operate, and, in particular, for network services including IPv6, multicast, Quality of Service (QoS), satellite-based communication, and security. Applications for this emulation program included communication of a variety of multimedia-based services, including video conferencing and video-on-demand.

26. In addition to having co-founded a technology company myself, I have worked for, consulted with, and collaborated with companies for nearly 30 years. These companies range from well-established companies to start-ups and include IBM, Hitachi Telecom, Turner Broadcasting System (TBS), Bell South, Digital Fountain, RealNetworks, Intel Research, Cisco Systems, and Lockheed Martin.

27. I am a Member of the Association of Computing Machinery (ACM) and a Fellow of the Institute of Electrical and Electronics Engineers (IEEE).

28. The billing rate for my services related to this matter is \$700 per hour. My compensation is in no way contingent on the outcome of this action.

29. Additional details about my employment history, fields of expertise, courses taught, and publications are further included in my Curriculum Vitae attached here as Appendix A.

III. LEGAL STANDARDS

30. I understand that claim construction begins with the language of the claims themselves. Claim terms are generally given their ordinary and customary meaning as understood by a person of ordinary skill in the art (“POSITA”) when viewing the claim terms in the context of the entire patent.

31. I understand that, in some cases, the plain and ordinary meaning of a claim term may be readily apparent and claim construction in such cases involves little more than the application of the widely accepted meaning of commonly understood words.

32. I understand that, in other cases, a claim term may have a specialized meaning in which case it is often necessary to look to the intrinsic evidence—which I understand to include the claims, the specification, and the prosecution/file history of the patent at issue—to construe the claim term. Indeed, I understand that the context in which a term is used in a claim can be highly instructive. I also understand that the specification is highly relevant to claim construction and can be the single best guide in determining the meaning of a claim term. In this respect, I understand that a claim construction that stays true to the claim language and most naturally aligns with the specification will be the correct construction. Accordingly, I understand that I must refrain from importing limitations into the claims that are not required by the intrinsic evidence.

33. Further, I understand that extrinsic evidence – dictionaries, treatises, and the like – can also be used to assist with claim construction. However, I understand that intrinsic evidence is often more reliable than the extrinsic evidence.

IV. LEVEL OF ORDINARY SKILL IN THE ART

34. I have been asked to offer my opinion regarding the level of ordinary skill in the art with respect to the Zone Scene Patents.

35. To assess the level of ordinary skill in the art, I understand one considers the type of problems encountered in the art, the prior solutions to those problems, the rapidity with which innovations are made, the sophistication of the technology, and the level of education of active workers in the field.

36. To assess the level of ordinary skill in the art of the Zone Scene Patents here, I have reviewed the Zone Scene Patents and related documents and considered the type of problems encountered in the art, the prior solutions to those problems, the rapidity with which innovations are made, the sophistication of the technology, and the level of education of active workers in the field. In addition, I considered my own experience teaching and performing research in the networking and consumer audio systems fields, as well as my experience collaborating and consulting with concerns in these industries.

37. I understand that neither party has taken a position yet regarding the level of ordinary skill in the art for the asserted patents. Based on my assessment, it is my opinion that a POSITA for purposes of this action is a person having the equivalent of a four-year degree from an accredited institution (typically denoted as a B.S. degree) in computer science, computer engineering, electrical engineering, or an equivalent thereof, and approximately 2-4 years of

professional experience in the fields of networking and network-based systems or applications, such as consumer audio systems, or an equivalent level of skill, knowledge, and experience.

38. Based on my aforementioned analysis and my personal knowledge and experience in the fields of networking and consumer audio systems, including the configuration and/or control of networked devices, I am very familiar with the level of knowledge and abilities of a POSITA at the time of the inventions.

39. In forming the opinions set forth herein, I applied this level of ordinary skill in the art.

V. BRIEF OVERVIEW OF THE ZONE SCENE PATENTS

40. As noted above, the Zone Scene Patents are part of the same patent family and claim priority to the ‘407 Provisional, which was filed on September 12, 2006. Further, the ‘966 and ‘855 patents are both direct continuations of U.S. Non-Provisional Application No. 15/130,919 (the “‘919 Application”), which was filed on April 15, 2016, and the ‘919 Application is a direct continuation of the ‘206 patent. In this Declaration, my references below to the teachings of the Zone Scene Patents are based on the column and line numbers of the ‘206 patent’s specification, but it should be understood that the same such teachings are likewise found in the ‘966 and ‘885 patent’s specification.

41. The ‘206 patent discloses a “multi-zone system” (also referred to as an “entertainment system”) comprising one or more “zone players” (also referred to as “players” or “playback devices”) that are coupled to a “data network” and are capable of playing multimedia, such as audio, in “digital” format. *See, e.g.*, ‘206 Patent at 2:28-37; 3:13-15; 4:31-5:3; 5:9-6:5, FIG. 1. The ‘206 patent illustrates an example system configuration as follows:

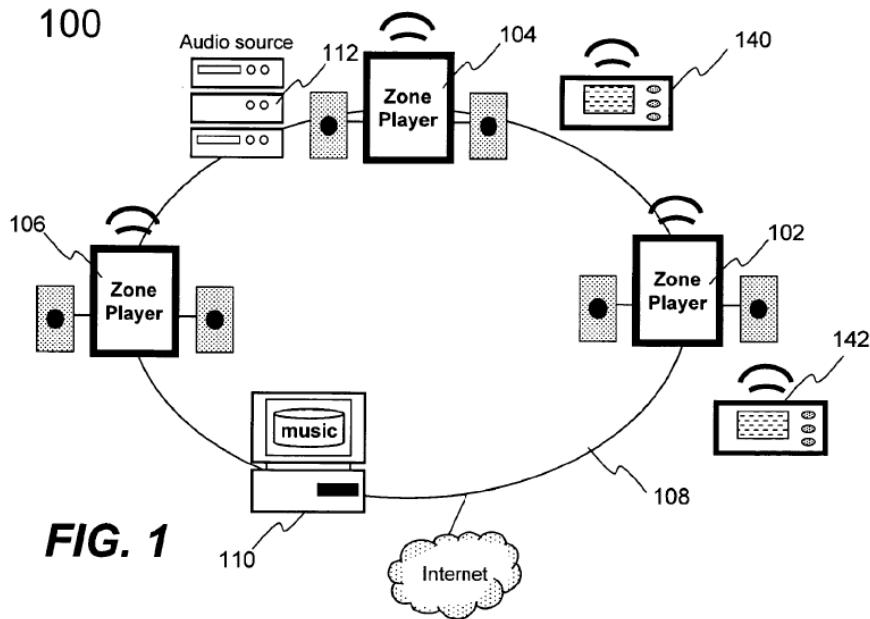


FIG. 1

42. In this configuration, each “zone player” (e.g., zone player 102, 104, or 106) is configured to communicate over a “data network” (e.g., data network 108, which may take the form of a “local area network”) with various other devices, including one or more other “zone players,” one or more “controlling devices” (e.g., controlling device 140 or 142 that are also referred to as a “controller”), and one or more local audio sources. *See, e.g.*, ‘206 Patent at 4:26-5:8, 6:12-31, FIG. 1. Additionally, each “zone player” is also configured to communicate over a “wide area network” (via the “local area network”) with one or more remote audio sources (e.g., an Internet-based audio source). *See, e.g., id.* at 4:54-64, FIG. 1.

43. As disclosed in the ‘206 patent, a mechanism is provided to place “zone players” into a “zone scene,” which is a predefined grouping of “zone players” that can first be saved by a user and can then later be invoked in order to cause the predefined grouping of “zone players” to become configured for synchronous playback of media. *See, e.g., id.* at Abstract, 3:5-21; 8:24-36. Thus, according to the ‘206 Patent, “zone players” may communicate over the “data network” for purposes of synchronized playback of media retrieved from an audio source, among

other purposes. *See, e.g., id.* at Abstract, 3:5-21; 4:37-39; 7:16-19; 7:40-45; FIG. 1. The ‘206 patent states that “unless explicitly stated otherwise, an audio source or audio sources are in digital format and can be transported or streamed over a data network.” *Id.* at 4:37-39. Further, the ‘206 Patent explains that communications over the “data network” are in the form of digital data “packets” and are in accordance with one or more common communication protocols, such as IEEE 802.11a, 802.11b, 802.11g. *See, e.g., id.* at 4:50-55; 5:14-39; 6:21-23.

VI. ASSERTED CLAIMS OF THE ZONE SCENE PATENTS

44. I understand that the term “data network” appears in dependent claims 7 and 15 of the ‘966 patent and all asserted claims of the ‘885 patent. I have reproduced representative claims 1 and 7 of the ‘966 patent and claim 1 of the ‘885 patent below:

‘966 patent, claim 1

[1.0] A *computing device* comprising:

[1.1] one or more processors;

[1.2] a non-transitory *computer-readable medium*; and

[1.3] program instructions stored on the non-transitory computer-readable medium that, when executed by the one or more processors, cause the computing device to perform functions comprising:

[1.4] while serving as a controller for a networked media playback system comprising a first zone player and at least two other zone players, wherein the first zone player is operating in a standalone mode in which the first zone player is configured to play back media individually:

[1.5] receiving a first request to create a first zone scene comprising a first predefined grouping of zone players including at least the first zone player and a second zone player that are to be configured for synchronous playback of media when the first zone scene is invoked;

[1.6] based on the first request, i) causing creation of the first zone scene, ii) causing an indication of the first zone scene to be transmitted to the first zone player, and iii) causing storage of the first zone scene;

[1.7] receiving a second request to create a second zone scene comprising a second predefined grouping of zone players including at least the first zone player and a third zone player that are to be configured for synchronous playback of media when the second zone scene is invoked, wherein the third zone player is different than the second zone player;

[1.8] based on the second request, i) causing creation of the second zone scene, ii) causing an indication of the second zone scene to be transmitted to the first zone player, and iii) causing storage of the second zone scene;

[1.9] displaying a representation of the first zone scene and a representation of the second zone scene; and

[1.10] while displaying the representation of the first zone scene and the representation of the second zone scene, receiving a third request to invoke the first zone scene; and

[1.11] based on the third request, causing the first zone player to transition from operating in the standalone mode to operating in accordance with the first predefined grouping of zone players such that the first zone player is configured to coordinate with at least the second zone player to output media in synchrony with output of media by at least the second zone player.

‘966 patent, claim 7

[7.0] The *computing device* of claim 1, further comprising program instructions stored on the non-transitory *computer-readable medium* that, when executed by the one or more processors, cause the computing device to perform functions comprising:

before displaying the representation of the first zone scene and the representation of the second zone scene, *receiving, from another device over a data network*, data defining the first zone scene and data defining the second zone scene.

‘885 patent, claim 1

[1.0] A *first zone player* comprising:

[1.1] a network interface that is configured to communicatively couple the first zone player to at least one *data network*;

[1.2] one or more processors;

[1.3] a non-transitory *computer-readable medium*; and

[1.4] program instructions stored on the non-transitory computer-readable medium that, when executed by the one or more processors, cause the first zone player to perform functions comprising:

[1.5] while operating in a standalone mode in which the first zone player is configured to play back media individually in a networked media playback system comprising the first zone player and at least two other zone players:

[1.6] (i) *receiving, from a network device over a data network*, a first indication that the first zone player has been added to a first zone scene comprising a first predefined grouping of zone players including at least the first zone player and a second zone player that are to be configured for synchronous playback of media when the first zone scene is invoked; and

[1.7] (ii) *receiving, from the network device over the data network*, a second indication that the first zone player has been added to a second zone scene comprising a second predefined grouping of zone players including at least the first zone player and a third zone player that are to be configured for synchronous playback of media when the second zone scene is invoked, wherein the second zone player is different than the third zone player;

[1.8] after receiving the first and second indications, continuing to operate in the standalone mode until a given one of the first and second zone scenes has been selected for invocation;

[1.9] after the given one of the first and second zone scenes has been selected for invocation, *receiving, from the network device over the data network*, an instruction to operate in accordance with a given one of the first and second zone scenes respectively comprising a given one of the first and second predefined groupings of zone players; and

[1.10] based on the instruction, transitioning from operating in the standalone mode to operating in accordance with the given one of the first and second predefined groupings of zone players such that the first zone player is configured to coordinate with at least one other zone player in the given one of the first and second predefined groupings of zone players *over a data network* in order to output media in synchrony with output of media by the at least one other zone player in the given one of the first and second predefined groupings of zone players.

VII. BACKGROUND ON “DATA NETWORK”

45. In the field of networking, the term “data network” is a well-understood term of art that is often used interchangeably with other well-understood terms such as “computer network,” “packet network,” and “data communications network” to refer to a specific class of networks that interconnect and enable devices to exchange information with one another in the form of digital data packets. *See, e.g.*, Appendix B (SONOS-SVG2-00018673-700) at 676, Appendix C (SONOS-SVG2-00018417-22) at 22; Appendix D (SONOS-SVG2-00018301-13) at

305; Appendix E (SONOS-SVG2-00018832-36) at 36. In this respect, a POSITA would understand that a “data network” has at least two defining characteristics that distinguish it from other types of communication mediums that are not included in this class of networks.

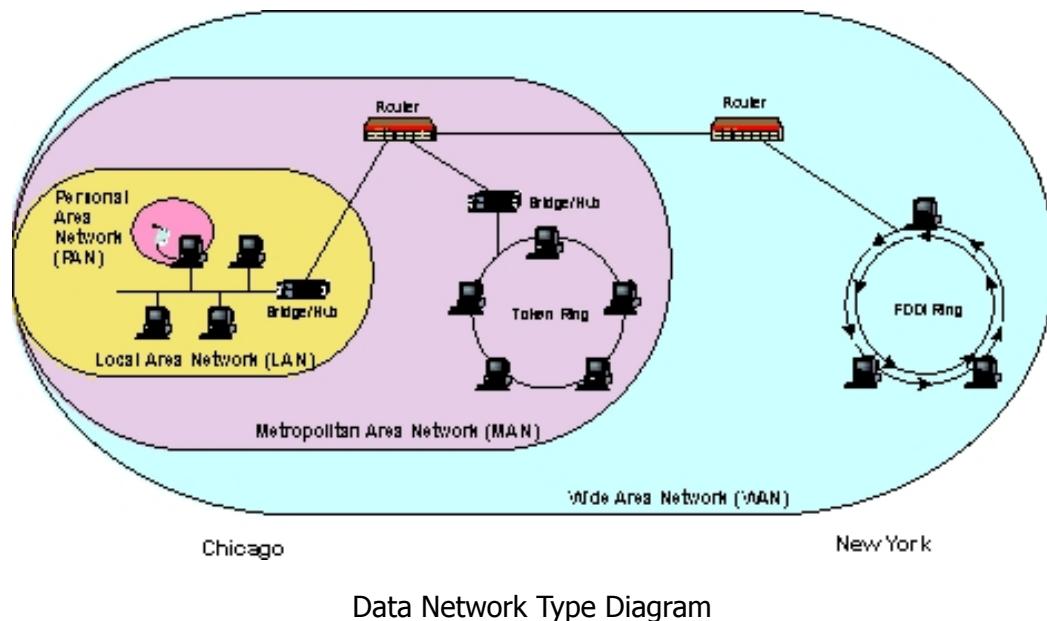
46. First, a POSITA would understand that a “data network” enables networked devices to engage in ***two-way*** communication with one another, which distinguishes it from a medium that only allows for one-way communication. In other words, a “data network” is commonly understood to interconnect network devices together (*e.g.*, a computers) such that a given network device can both send information to other network devices on the “data network” and also receive information from other network devices on the “data network.”

47. Second, a POSITA would understand that a “data network” transfers information between networked devices in the form of ***digital data packets*** (sometimes more generally referred to as “computer data”), which distinguishes it from a medium that transfers other forms of information. As understood in the field of networking, a digital data packet is a unit of digital data in a “data network.” In practice, when information is sent from one networked device to another over a “data network,” the information is broken down into units of digital data called “packets” (sometimes also referred to as “frames”), and these digital data packets are sent individually and reassembled once they reach another interconnected device (*i.e.*, the destination) on the “data network.”

48. The digital data packets that are exchanged over a “data network” will have a specific structure, which is typically defined by the communications protocol that is utilized by the “data network.” For instance, Transmission Control Protocol/Internet Protocol (“TCP/IP”), which is a common Internet protocol used to transfer packets over a “data networks,” describes that a digital data packet comprises a header and a payload.

49. The header of a digital data packet generally includes information about the digital data packet, such as the length of the packet, the packet number (e.g., a number that identifies which packet this is in a sequence of packets), the communications protocol, the source address (e.g., an IP address that corresponds to where the packet came from), and the destination address (e.g., an IP address that corresponds to where the packet is going). In turn, the payload (also referred to as “data”) of a digital data packet generally includes the actual information that the packet is delivering to the destination.

50. As well understood in the field of networking, and as illustrated below, “data networks” can also be categorized into types depending on their coverage range. Some common examples of “data network” types are shown in the following figure that is found in the online *VOIP Industry Dictionary*:



Appendix F (SONOS-SVG2-00018823-24) at 23 (describing and showing the “basic types of data networks”).

51. As illustrated, one type of “data network” that is well understood in the field of

networking is known as a “local area network” (“LAN”), which generally covers a relatively small geographic area, such as a house or a building. A LAN can be implemented via various protocols, which are standards that define how information is exchanged between interconnected devices. Example LAN protocols may include Ethernet, token ring, Asynchronous Transfer Mode (“ATM”), Fiber Distributed Data Interface (“FDDI”) II, and 10Base-T. Ethernet is the most common type of LAN technology and is recognizable by its common “CAT5” copper cable medium.

52. It is also well understood in the field of networking that a LAN may take the form of a Wireless LAN (“WLAN”), which is commonly referred to as “Wi-Fi.” A WLAN operates essentially the same as a traditional LAN, but uses wireless signals between antennas as the medium, rather than cabling. Similar to a LAN, a WLAN may be implemented via various protocols. Most of the protocols used at the time of the claimed inventions and still in use today are based on the 802.11 standard and are differentiated by the letter appearing after the number. For instance, the three main protocols used at the time of the claimed inventions are 802.11a, 802.11b, and 802.11g.

53. Another type of “data network” that is well understood in the field of networking is known as a wide area network (“WAN”) that generally spans longer distances than a LAN. A WAN often serves as an interconnection between multiple LANs. The Internet is a well-known example of a WAN. A WAN may also be implemented via various protocols, examples of which may include Frame Relay, X.25, Integrated Services Digital Network (“ISDN”), and Point-to-Point Protocol (“PPP.”).

54. In summary, a POSITA would understand that a “data network” (which may also be referred to as a “computer network,” “data communications network,” or “packet network”)

enables interconnected devices to send digital data packets to and receive digital data packets from each other.

VIII. PROPOSED CONSTRUCTIONS OF “DATA NETWORK”

55. Each party’s proposed construction is set forth in the following table.

Sonos’s Proposed Construction	Google’s Proposed Construction
“a medium that interconnects devices, enabling them to send digital data packets to and receive digital data packets from each other”	Plain and ordinary meaning; no construction necessary at this time

56. It is my opinion that Sonos’s construction is entirely consistent with how a POSITA would understand the term “data network” in the context of the Zone Scene Patents because that construction properly specifies that a “data network” requires (i) *two-way* communication of information (*i.e.*, sending and receiving information) that (ii) exchange of information in the form of *digital data packets*.

57. While Google has asserted the “plain and ordinary meaning” of this term, I understand that Google disagrees with Sonos’s proposed construction and has not explained what this plain and ordinary meaning of “data network” should be. Thus, it is my opinion that Google’s construction of the term could improperly broaden the meaning of the term to cover other kinds of “networks” that do not constitute *data* networks. For instance, it is my opinion that Google’s construction creates a risk of improperly broadening the meaning of “data network” to cover (i) one-way communication of information, and/or (ii) exchange of information that may take various forms other than digital data packets.

58. Accordingly, Sonos has asked me to provide my opinions regarding how a POSITA would interpret the term “data network” in the context of the Zone Scene Patents. I have reached these opinions after reviewing the intrinsic evidence of the Zone Scene Patents and

performing an analysis of how a POSITA would understand the meaning of the term “data network” in the context of these patents. Additionally, I have also reviewed various technical sources explaining the term “data network.” The following paragraphs provide my explanation of the bases for my opinions set forth herein.

A. The Intrinsic Evidence Supports Sonos’s Proposed Construction

59. After reviewing the Zone Scene Patents, it is my opinion that the term “data network” is being used consistently with its well-understood meaning as set forth above, such that a POSITA would likewise understand the term “data network,” in the context of the Zone Scene Patents, requires *two-way* communication of information that takes the form of *digital data packets*.

60. For instance, the claim language of the ‘966 and ‘885 patents provides support for my opinion that the claimed “data network” requires *two-way* communication of information. Specifically, as reproduced above, claim 1 of the ‘885 patent expressly recites that the claimed “first *zone player*” is “communicatively couple[d] . . . to at least one data network” and is configured to *receive* data from a “network device” (e.g., a *controller*) “over a data network[.]” Similarly, dependent claim 7 of the ‘966 patent expressly recites that the claimed “computing device,” “while serving as a *controller* for a networked media playback system comprising a first zone player and at least two other zone players” as recited in claim 1, is configured to *receive* data from “another device” (e.g., one or more “*zone players*” in the “networked media playback system”) over a “data network.” In this regard, the claimed “computing device” (e.g., a controller) and the “first zone player” may be configured to send and receive data from each other.

61. The claim language of the ‘966 and ‘885 patents also provides support for my

opinion that the claimed “data network” requires information to be exchanged in the form of ***digital data packets***. For instance, the ‘966 patent recites a “***computing*** device” that comprises “***computer***-readable medium,” and the ‘885 patent recites a “first zone player” that comprises “***computer***-readable medium.” A POSITA would understand that these networked devices are thus computers, and computers that communicate over a “data network” exchange information in the form of digital data packets.

62. Consistent with the claim language, the specifications of the Zone Scene Patents repeatedly and uniformly describe that the claimed “data network” (e.g., data network 108) must enable ***two-way*** communication (i.e., sending and receiving information) between connected devices. For instance, the ‘206 patent discloses the following:

The network interface 202 facilitates ***a data flow between a data network (i.e., the data network 108 of FIG. 1) and the zone player 200*** and typically executes a special set of rules (i.e., a protocol) to ***send data back and forth***. One of the common protocols used in the Internet is TCP/IP (Transmission Control Protocol/Internet Protocol). In general, a network interface manages the assembling of an audio source or file into smaller ***packets that are transmitted over the data network*** or ***reassembles received packets*** into the original source or file. In addition, the network interface 202 handles the ***address part of each packet so that it gets to the right destination or intercepts packets destined for the zone player 200***.

‘206 Patent at 5:14-26;

The network interface 202 may include one or both of a wireless interface 216 and a wired interface 217. The wireless interface 216, also referred to as a RF interface, provides network interface functions by a ***wireless means for the zone player 200 to communicate with other devices*** in accordance with a communication protocol (such as the wireless standard IEEE 802.11a, 802.11b or 802.11g). The wired interface 217 provides network interface functions by a wired means (e.g., an Ethernet cable). In one embodiment, a zone player includes both of the interfaces 216 and 217, and other zone players include only a RF or wired interface. ***Thus these other zone players communicate with other devices on a network or retrieve audio sources via the zone player.***

Id. at 5:27-39;

According to one embodiment of the present invention, the memory 206 is used to save one or more saved zone configuration files that may be *retrieved* for modification at any time. Typically, a saved zone group configuration file is *transmitted* to a controller (e.g., the controlling device 140 or 142 of FIG. 1, a computer, a portable device, or a TV) when a user operates the controlling device.

Id. at 5:51-57;

Many devices on the *[data] network 108* are configured to *download* and store audio sources. For example, the computing device 110 can *download audio sources* from the Internet and store the downloaded sources locally *for sharing with other devices on the Internet or the network 108*. The computing device 110 or any of the zone players can also be configured to *receive* streaming audio.

Id. at 4:58-63;

The controller 270 includes a network interface 280 referred to as a RF interface 280 that facilitates wireless communication with a zone player via a corresponding RF interface thereof. In one embodiment, the commands such as volume control and audio playback synchronization are sent via the RF interfaces. In another embodiment, a saved zone group configuration is *transmitted between a zone player and a controller* via the RF interfaces.

Id. at 7:26-33;

In one embodiment, a user creates a zone group including at least two zone players from *the controller 240 that sends signals or data to one of the zone players. As all the zone players are coupled on a network, the received signals in one zone player can cause other zone players in the group to be synchronized* so that all the zone players in the group playback an identical audio source or a list of identical audio sources in a timely synchronized manner.

Id. at 7:38-45; *see also, e.g., id.* at 4:37-39 (“As used herein, unless explicitly stated otherwise, an audio source or audio sources are in digital format and can be *transported or streamed over a data network.*”); 5:1-3 (“In accordance with the present invention, the audio source may be *shared among the devices* on the *[data] network 108.*”); 10:27-30 (“[T]he interconnections of the players are checked to make sure that the players *communicate among themselves* and/or with a controller if there is such a controller in the scene.”); 3:13-15;4:54-57; 6:2-5, 6:16-31; 10:40-45, FIG. 1.

63. The Zone Scene Patents also explain that information that is exchanged over a “data network” takes the form of **digital data packets**, which is consistent with how a POSITA would understand the term. *See, e.g.*, ‘206 Patent at 5:14-26:

The network interface 202 facilitates a ***data flow*** between ***a data network (i.e., the data network 108 of FIG. 1)*** and the zone player 200 and typically executes a special set of rules (i.e., a protocol) to send data back and forth. One of the common protocols used in the Internet is TCP/IP (Transmission Control Protocol/Internet Protocol). In general, a network interface manages the assembling of an audio source or file into smaller ***packets*** that are transmitted over the data network or reassembles received ***packets*** into the original source or file. In addition, the network interface 202 handles the ***address part of each packet*** so that it gets to the right destination or intercepts ***packets*** destined for the zone player 200.

‘206 Patent at 5:14-26.

64. The Zone Scene Patents also repeatedly and uniformly explain that information that is exchanged over the “data network” is in a “***digital***” format. *See, e.g., id.* at 4:37-39 (“As used herein, ***unless explicitly stated otherwise***, an audio source or audio sources are in ***digital*** format and can be transported or streamed over a ***data network**digital*** format and retrieved from a source over a network.”); 6:2-5 (“[T]he audio processing circuit 210 may include necessary circuitry to process analog signals as inputs to produce ***digital*** signals for sharing with other devices on a network.”).

65. Further, the Zone Scene Patents repeatedly and uniformly explain that the communications over a “data network” are performed in accordance with communication protocols that a POSITA would understand to involve digital data packets. *See, e.g., id.* at *Id.* at 5:14-20 (disclosing that “[t]he network interface 202 facilitates a data flow between a data network (i.e., the data network 108 of FIG. 1) and the zone player 200 and typically executes a special set of rules (i.e., a ***protocol***) to send data back and forth,” and “[o]ne of the common

protocols used in the Internet is ***TCP/IP (Transmission Control Protocol/Internet Protocol).***”); 4:41-54 (“[A]ll devices including the zone players 102, 104 and 106 are coupled to the network 108 by wireless means ***based on an industry standard such as IEEE 802.11.***”); 5:28-33 (“The wireless interface 216 … provides network interface functions by a wireless means for the zone player 200 to communicate with other devices ***in accordance with a communication protocol (such as the wireless standard IEEE 802.11a, 802.11b or 802.11g).***”)

66. Thus, it is my opinion that this repeated and uniform disclosure would lead a POSITA to understand that the term “data network,” in the context of the Zone Scene Patents, refers to a medium that interconnects devices, enabling them to ***send digital data packets to and receive digital data packets*** from each other.

B. The Extrinsic Evidence is Consistent with Sonos’s Proposed Construction

67. My opinion that a “data network,” in the context of the Zone Scene Patents, refers to a medium that interconnects devices, enabling them to send digital data packets to and receive digital data packets from each other is also confirmed by the various technical sources provided below.

68. As one example, *Data & Computer Communications* (6th ed. 2000) begins by explaining that the “two major categories into which communications networks are traditionally classified” are “wide area networks (WANs) and local area networks (LANs).” Appendix G (SONOS-SVG2-00018715-24) at 18. It goes on to explain that, “[a]s with WANs, a LAN is a communication network that interconnects a variety of devices and provides ***a means for information exchange*** among those devices” and that, “[a]t each station, there is a ***transmitter/receiver*** that communicates over a medium shared by other stations.” *Id.* at 21. This discussion confirms that, as explained above, a “data network” enables the interconnected

devices to engage in *two-way* communication.

69. *Data & Computer Communications* (6th ed. 2000) further explains that WANs (which as explained above are a common type of a “data network”) are generally implemented using “packet switching” and “ATM” technologies. *Id.* at 18-20. With respect to “packet switching” technology, it explains that “**data** are *sent out* in a sequence of small chunks, called **packets**. . . . Packet-switching networks are commonly used for terminal-to-computer and computer-to-computer communications.” *Id.* at 19. It also explains that “packet switching was developed at a time when **digital** long-distance transmission facilities exhibited a relatively high error rate compared to today’s facilities.” *Id.* With respect to “ATM” technology, it explains:

Asynchronous transfer mode (ATM), sometimes referred to as cell relay, is a culmination of all of the developments in circuit switching and packet switching over the past 25 years. ATM can be viewed as an evolution from frame relay. The most obvious difference between frame relay and ATM is that frame relay uses **variable-length packets, called frames**, and ATM uses **fixed-length packets, called cells**.

Id. at 20. This discussion confirms that, as explained above, a “data network” carries information in the form of ***digital data packets***.

70. As another example, *Encyclopedia of Networking* (1998) explains that, in a LAN (which as explained above are a common type of a “data network”), “once a workstation is ready to transmit and has access to the **shared medium**, it simply puts the **packets** on the network and hopes that the recipient receives them.” Appendix H (SONOS-SVG2-00018707-14) at 10. It also provides the following explanation of how information is packaged:

Data is packaged into **frames** for transmission on the LAN. . . . A **frame** is usually **addressed** for a single computer, although a multicast **address** can be used to transmit to all workstations on the LAN. Higher-layer protocols such as **IP** and **IPX** package **data into datagrams**. **Datagrams** are in turn divided up and put into **frames** for transmission on a particular LAN.

Id. This discussion likewise confirms that a “data network” enables *two-way* communication

(which is discussed here in the context of a “shared medium”) and carries information in the form of ***digital data packets*** (referred to here as “frames”).

71. As yet another example, *Computer Networks and Internets* (2nd ed. 1999) explains that “[e]ach LAN consists of a single ***shared medium***” and “computers take turns using the medium to send ***packets***.” Appendix D (SONOS-SVG2-00018301-13) at 05. It goes on to explain that “computers take turns using the medium to send data” and “[a]lthough ***LAN technologies require computers to divide data into small packets called frames***, only one ***packet*** can be transmitted on a LAN at any time.” *Id.* at 12. As with the preceding example, this discussion confirms that a “data network” enables ***two-way*** communication (which is again discussed here in the context of a “shared medium”) and carries information in the form of ***digital data packets*** (referred to here as both “packets” and “frames”).

72. Various other definitions of a LAN likewise demonstrate that a “data network” has these characteristics.

73. For example, *Data Communications and Computer Networks for Computer Scientists and Engineers* (2nd ed. 2003) defines a LAN as “[a] data communications network used to ***interconnect*** a community of ***digital*** devices over a localized area” and “[m]essages within a LAN are transmitted as a series of variable length ***frames*** using transmission media which introduce only relatively low error rates.” Appendix C (SONOS-SVG2-00018417-22) at 20, 22. It also provides some history explaining that “[m]any modern LANs evolved from a LAN known as Aloha which was . . . ***packet*** based and used radio as its transmission medium.” *Id.* at 21.

74. As another example, *Encyclopedia of Computer Science and Technology* (2009) provides the following explanation regarding a LAN:

There are two basic ways to connect computers in a LAN. The first [is] called Ethernet . . . Ethernet uses a single cable line called a bus to which all participating computers are connected. Each **data packet** is received by all computers, but processed only by the one it is addressed to. . . . Naturally there must be software to manage the **transmission and reception** of **data packets**. The structure of a **packet** (*sometimes called a frame*) has been standardized

Appendix I (SONOS-SVG2-00018402-06) at 06. This discussion confirms that a “data network” enables **two-way** communication (*i.e.*, “transmission and reception”) and carries information in the form of **digital data packets**.

75. As yet another example, the *Webster’s New World Telecom Dictionary* (2008) defines “LAN” as “a **packet network** designed to **interconnect** host computers, peripherals, storage devices, and other computing resources within a local area, *i.e.*, limited distance.” Appendix E (SONOS-SVG2-00018832-36) at 36.

76. As another example, *Packet Broadband Network Handbook* (2004) provides the following explanation regarding a LAN:

A local area network is a high-speed **data network** that covers a relatively small geographic area. . . . LAN is a type of broadband **packet** access **network** that carries the **packet data** traffic of an organization. LAN interconnects the end users of an organization to an outside public data network such as the Internet. . . . The physical layer [of the LAN protocols] is primarily concerned with the **transmission medium** and its physical characteristics for **digital signal transmission**.⁷

Appendix B (SONOS-SVG2-00018673-700) at 676. This discussion confirms that a “data network” carries information in the form of **digital data packets**.

77. As a further example, *Local & Metropolitan Area Networks* (6th ed. 2000) states that “[a] communications network is a facility that interconnects a number of devices and provides a **means for transmitting data from one** attached device **to another**.⁸” Appendix J (SONOS-SVG2-00018752-64) at 57. It further explains that each device that attaches to a LAN typically has a “network interface card (NIC)” that contains logic “**for sending and receiving**

blocks of data on the LAN.” *Id.* at 59. This discussion confirms that, as explained above, a “data network” enables the interconnected devices to engage in *two-way* communication.

78. Thus, consistent with the intrinsic evidence, the extrinsic evidence overwhelmingly support my opinion that a POSITA would understand the term “data network” in the context of the Zone Scene Patents to require *two-way* communication of information in the form of *digital data packets*.

IX. CONCLUSION

79. In view of the foregoing, it is my opinion that Sonos’s proposed claim construction for “data network” accurately captures how a POSITA should interpret that term in the context of the Zone Scene Patents.

I declare under penalty of perjury that the foregoing is true and correct to the best of my knowledge and belief

Dated: April 27, 2021

By: Kevin C. Almeroth
Kevin C. Almeroth

Appendix A

Kevin C. Almeroth

Professor Emeritus
 Department of Computer Science
 University of California
 Santa Barbara, CA 93106-5110
 (805)636-1123 (office)
 (805)893-8553 (fax)
 almeroth@cs.ucsb.edu (email)
<http://www.cs.ucsb.edu/~almeroth> (WWW URL)

Education

Ph.D.	June 1997 <i>Georgia Institute of Technology</i> Computer Science <i>Dissertation Title:</i> Networking and System Support for the Efficient, Scalable Delivery of Services in Interactive Multimedia Systems <i>Minor:</i> Telecommunications Public Policy
M.S.	June 1994 <i>Georgia Institute of Technology</i> Computer Science <i>Specialization:</i> Networking and Systems
B.S. (high honors)	June 1992 <i>Georgia Institute of Technology</i> Information and Computer Science <i>Minors:</i> Economics, Technical Communication, American Literature

Employment History

Professor Emeritus	Department of Computer Science University of California Santa Barbara, CA	Nov 2020 -- present
Professor	Department of Computer Science University of California Santa Barbara, CA	Jul 2005 -- Oct 2020
Associate Dean	College of Engineering University of California Santa Barbara, CA	Mar 2007 -- Aug 2009
Vice Chair	Department of Computer Science University of California Santa Barbara, CA	Jul 2000 -- Nov 2005

Associate Professor	Department of Computer Science University of California Santa Barbara, CA	Jul 2001 -- Jun 2005
Assistant Professor	Department of Computer Science University of California Santa Barbara, CA	Jul 1997 -- Jun 2001
Graduate Researcher	Broadband Telecommunications Center Georgia Center for Adv Telecom Tech Atlanta, GA	Sep 1996--Jun 1997
Graduate Intern	IBM T.J. Watson Research Labs Hawthorne, NY	Jun 1995--Sep 1995
Support Specialist	Office of Information Technology Georgia Institute of Technology Atlanta, GA	Sep 1995--Jun 1997
Research Assistant	College of Computing Georgia Institute of Technology Atlanta, GA	Jan 1994--Mar 1994
Graduate Intern	Hitachi Telecommunications Norcross, GA	Jun 1992--Sep 1992
Undergraduate Intern	IBM Research Triangle Park, NC	Jun 1989--Sep 1989 Jun 1990--Sep 1990 Mar 1991--Sep 1991

Industry Technical Advising

Board of Directors	<u>The New Media Studio</u> Santa Barbara, CA	Nov 2006 -- present
Co-Founder & Chairman of the Board	Santa Barbara Labs, LLC Santa Barbara, CA	Sep 2007 -- Dec 2009
Board of Advisors	Techknowledge Point Santa Barbara, CA	May 2001 -- Dec 2007
Technical Advisory Board	Occam Networks, Inc. Santa Barbara, CA	May 2000 -- Dec 2010
Board of Advisors	Airplay Inc. San Francisco, CA	Jun 2005 -- Aug 2009
Consultant	Lockheed Martin Corporation San Jose, CA	Nov 1999 -- Jun 2009

Board of Advisors	Santa Barbara Technology Group Santa Barbara, CA	Sep 2000 -- Dec 2004
Board of Directors	Virtual Bandwidth, Inc. Santa Barbara, CA	Nov 2000 -- Jun 2001
Board of Advisors & Affiliated Scientist	Digital Fountain San Francisco, CA	Jan 2000 -- Dec 2001
Senior Technologist	IP Multicast Initiative, Stardust Forums Campbell, CA	Jun 1998 -- Dec 2000

I. Teaching

A. Courses Taught

CS 176A	Intro to Computer Communication Networks	Fall 1997, Fall 1998, Fall 2002, Fall 2003, Fall 2004, Spring 2005, Spring 2006, Spring 2007, Spring 2008, Fall 2008, Fall 2009, Fall 2010, Fall 2011, Fall 2012, Fall 2013, Fall 2014, Spring 2017, Spring 2018, Spring 2020, Fall 2020
CS 176B	Network Computing	Winter 2000, Winter 2001, Winter 2002, Winter 2012, Winter 2014, Winter 2015, Winter 2018, Winter 2019, Winter 2020
MAT 201B	Media Networks and Services	Fall 1999, Fall 2000, Fall 2001, Fall 2003
CS 276	Distributed Computing and Computer Networks	Winter 1999, Spring 2000, Fall 2002, Fall 2005, Fall 2018
CS 290I	Networking for Multimedia Systems	Winter 1998, Spring 1999, Fall 2004, Winter 2010
CS 595N	Technology and Society	Winter 2005, Fall 2005, Spring 2006, Fall 2006, Spring 2007, Fall 2007, Spring 2008, Fall 2008, Spring 2009
CS 595N	Economic Systems Seminar	Winter 2004, Spring 2004, Winter 2005, Spring 2005
CS 595N	Networking Seminar	Winter 1999, Fall 1999, Winter 2003, Winter 2019
CS 595N	Wireless Networking & Multimedia Seminar	Fall 2000
CS 595I	Systems Design and Implementation Seminar	Fall 1999, Fall 2000, Winter 2001, Spring 2001, Winter 2002, Spring 2002

B. Other Teaching Experience

- *The Evolution of Advanced Networking Services: From the ARPAnet to Internet2*, Instructor, Summer 2001. Short course taught at Escuela de Ciencias Informatica (ECI) sponsored by the Universidad de Buenos Aires.

- *Johns Hopkins Center for Talented Youth*, Instructor, Summer 1994. CTY is a program to teach gifted high school students the fundamentals of computer science.
- *Georgia Institute of Technology*, Graduate Teaching Assistant, Sep 1994--Sep 1996. Worked as a TA for 12 quarters teaching 7 different courses (4 undergraduate and 3 graduate).

C. Ph.D. Students Advised [14 graduated]

14. Daniel Havey

Research Area: *Throughput and Delay on the Packet Switched Internet*

Date Graduated: Winter 2015

First Position: Microsoft

13. Lara Deek (co-advised with E. Belding)

Research Area: *Resource-Efficient Wireless Systems for Emerging Wireless Networks*

Date Graduated: Summer 2014

First Position: Post Doc, UIUC

12. Mike Wittie

Research Area: *Towards Sustained Scalability of Communication Networks*

Date Graduated: Summer 2011

First Position: Assistant Professor, Montana State University

11. Allan Knight

Research Area: *Supporting Integration of Educational Technologies and Research of Their Effects on Learning*

Date Graduated: Summer 2009

First Position: Research Scientist, Citrix Online

10. Hangjin Zhang

Research Area: *Towards Blended Learning: Educational Technology to Improve and Assess Teaching and Learning*

Date Graduated: Spring 2009

First Position: Microsoft

9. Gayatri Swamynathan

Dissertation Title: *Towards Reliable Reputations for Distributed Applications*

Date Graduated: Spring 2008

First Position: Zynga

8. Amit Jardosh (co-advised with E. Belding)

Dissertation Title: *Adaptive Large-Scale Wireless Networks: Measurements, Protocol Designs, and Simulation Studies*

Date Graduated: Fall 2007

First Position: Yahoo!

7. Khaled Harras

Dissertation Title: *Protocol and Architectural Challenges in Delay and Disruption Tolerant Networks*

Date Graduated: Summer 2007

First Position: Assistant Professor, Carnegie Mellon University

6. Krishna Ramachandran (co-advised with E. Belding)

Dissertation Title: *Design, Deployment, and Management of High-Capacity Wireless Mesh Networks*

Date Graduated: Winter 2006

First Position: Research Scientist, Citrix Online

5. Robert Chalmers

Dissertation Title: *Improving Device Mobility with Intelligence at the Network Edge*
 Date Graduated: Summer 2004
 First Position: President and CEO, Limbo.net

4. Prashant Rajvaidya

Dissertation Title: *Achieving Robust and Secure Deployment of Multicast*
 Date Graduated: Spring 2004
 First Position: President and CTO, Mosaic Networking

3. Sami Rollins

Dissertation Title: *Overcoming Resource Constraints to Enable Content Exchange Applications in Next-Generation Environments*
 Date Graduated: Spring 2003
 First Position: Assistant Professor, Mount Holyoke College

2. Srinivasan Jagannathan

Dissertation Title: *Multicast Tree-Based Congestion Control and Topology Management*
 Date Graduated: Spring 2003
 First Position: Consultant, Kelly & Associates

1. Kamil Sarac

Dissertation Title: *Supporting a Robust Multicast Service in the Global Infrastructure*
 Date Graduated: Spring 2002
 First Position: Assistant Professor, UT-Dallas

D. M.S. Students Advised (Thesis/Project Option) [19 graduated]

19. Neer Shey

Research Area: *Analyzing Content Distribution Through Opportunistic Contact for Smart Cellular Phones*
 Date Graduated: Spring 2010

18. Camilla Fiorese

Research Area: *Analysis of a Pure Rate-Based Congestion Control Algorithm*
 Date Graduated: Summer 2009

17. Brian Weiner

Research Area: *Multi-Socket TCP: A Simple Approach to Improve Performance of Real-Time Applications over TCP*
 Date Graduated: Fall 2007

16. Avijit Sen Mazumder

Research Area: *Facilitating Robust Multicast Group Management*
 Date Graduated: Fall 2005

15. Rishi Matthew

Thesis Title: *Providing Seamless Access to Multimedia Content on Heterogeneous Platforms*
 Date Graduated: Summer 2004

14. Camden Ho

Research Area: *Tools and Techniques for Wireless Network Management*
 Date Graduated: Spring 2004

13. Amit Jardosh (co-advised with E. Belding)

Research Area: *Realistic Environment Models for Mobile Network Evaluation*
 Date Graduated: Spring 2004

12. Nitin Solanki

Research Area: *SongWand: A Wireless Barcode Scanner Using Bluetooth Technology*

Date Graduated: Winter 2004

11. Vrishali Wagle (co-advised with E. Belding)
 Research Area: *An Ontology-Based Service Discovery Mechanism*
 Date Graduated: Winter 2004

10. Uday Mohan
 Thesis Title: *Scalable Service Discovery in Mobile Ad hoc Networks*
 Date Graduated: Spring 2003

9. Krishna Ramachandran
 Thesis Title: *Ubiquitous Multicast*
 Date Graduated: Spring 2003

8. John Slonaker
 Thesis Title: *Inductive Loop Signature Acquisition Techniques*
 Date Graduated: Spring 2002

7. Mohammad Battah
 Thesis Title: *Dedicated Short-Range Communications Intelligent Transportation Systems Protocol (DSRC-ITS)*
 Date Graduated: Spring 2002

6. Kevin Vogel
 Thesis Title: *Integrating E-Commerce Applications into Existing Business Infrastructures*
 Date Graduated: Spring 2001

5. Sami Rollins
 Thesis Title: *Audio XmL: Aural Interaction with XML Documents*
 Date Graduated: Winter 2000

4. Andy Davis
 Thesis Title: *Stream Scheduling for Data Servers in a Scalable Interactive TV System*
 Date Graduated: Spring 1999

3. David Makofske
 Thesis Title: *MHealth: A Real-Time Graphical Multicast Monitoring Tool*
 Date Graduated: Winter 1999

2. Prashant Rajvaidya
 Thesis Title: *MANTRA: Router-Based Monitoring and Analysis of Multicast Traffic*
 Date Graduated: Winter 1999

1. Alex DeCastro (co-advised with Yuan-Fang Wang)
 Thesis Title: *Web-Based Collaborative 3D Modeling*
 Date Graduated: Winter 1998

E. Teaching Awards

2006-2007 UCSB Academic Senate Distinguished Teaching Award
 2004-2005 Computer Science Outstanding Faculty Member
 2000-2001 UCSB Spotlight on Excellence Award
 1999-2000 Computer Science Outstanding Faculty Member (co-recipient)
 1998-1999 Computer Science Outstanding Faculty Member (co-recipient)
 1997-1998 Computer Science Outstanding Faculty Member

II. Research

A. Journal Papers, Magazine Articles, Books, and Book Chapters

62. L. Deek, E. Garcia-Villegas, E. Belding, S.J. Lee, and K. Almeroth, "[A Practical Framework for 802.11 MIMO Rate Adaptation](#)," [Computer Networks](#), vol. 83, num. 6, pp. 332-348, June 2015.
61. L. Deek, E. Garcia-Villegas, E. Belding, S.J. Lee, and K. Almeroth, "[Intelligent Channel Bonding in 802.11n WLANs](#)," [IEEE Transactions on Mobile Computing](#), vol. 13, num. 6, pp. 1242-1255, June 2014.
60. H. Zhang and K. Almeroth, "[Alternatives for Monitoring and Limiting Network Access to Students in Network-Connected Classrooms](#)," [Journal of Interactive Learning Research \(JILR\)](#), vol. 24, num. 3, pp. 237-265, July 2013.
59. M. Tavakolifard and K. Almeroth, "[A Taxonomy to Express Open Challenges in Trust and Reputation Systems](#)," [Journal of Communications](#), vol. 7, num. 7, pp. 538-551, July 2012.
58. M. Tavakolifard and K. Almeroth, "[Social Computing: An Intersection of Recommender Systems, Trust/Reputation Systems, and Social Networks](#)," [IEEE Network](#), vol. 26, num. 4, pp. 53-58, July/August 2012.
57. M. Tavakolifard, K. Almeroth, and P. Ozturk, "[Subjectivity Handling of Ratings for Trust and Reputation Systems: An Abductive Reasoning Approach](#)," [International Journal of Digital Content Technology and its Applications \(JDCTA\)](#), vol. 5, num. 11, pp. 359-377, November 2011.
56. R. Raghavendra, P. Acharya, E. Belding and K. Almeroth, "[MeshMon: A Multi-Tiered Framework for Wireless Mesh Network Monitoring](#)," [Wireless Communications and Mobile Computing \(WCMC\) Journal](#), vol. 11, num. 8, pp. 1182-1196, August 2011.
55. A. Knight and K. Almeroth, "[Automatic Plagiarism Detection with PAIRwise 2.0](#)," [Journal of Interactive Learning Research \(JILR\)](#), vol. 22, num. 3, pp. 379-400, July 2011.
54. V. Kone, M. Zheleva, M. Wittie, B. Zhao, E. Belding, H. Zheng, and K. Almeroth, "[AirLab: Consistency, Fidelity and Privacy in Wireless Measurements](#)," [ACM Computer Communications Review](#), vol. 41, num. 1, pp. 60-65, January 2011.
53. G. Swamynathan, K. Almeroth, and B. Zhao, "[The Design of a Reliable Reputation System](#)," [Electronic Commerce Research Journal](#), vol. 10, num. 3-4, pp. 239-270, December 2010.
52. P. Acharya, A. Sharma, E. Belding, K. Almeroth and K. Papagiannaki, "[Rate Adaptation in Congested Wireless Networks through Real-Time Measurements](#)," [IEEE Transactions on Mobile Computing](#), vol. 9, num. 11, pp. 1535-1550, November 2010.
51. R. Raghavendra, E. Belding, K. Papagiannaki, and K. Almeroth, "[Unwanted Link Layer Traffic in Large IEEE 802.11 Wireless Networks](#)," [IEEE Transactions on Mobile Computing](#), vol. 9, num. 9, pp. 1212-1225, September 2010.
50. H. Zhang and K. Almeroth, "[Moodog: Tracking Student Activity in Online Course Management Systems](#)," [Journal of Interactive Learning Research \(JILR\)](#), vol. 21, num. 3, pp. 407-429, July 2010.
49. R. Chertov and K. Almeroth, "[Qualitative Comparison of Link Shaping Techniques](#)," [International](#)

Journal of Communication Networks and Distributed Systems, vol. 5, num. 1/2, pp. 109-129, July 2010.

48. A. Knight and K. Almeroth, "[Fast Caption Alignment for Automatic Indexing of Audio](#)," International Journal of Multimedia Data Engineering and Management, vol. 1, num. 2, pp. 1-17, April-June 2010.
47. K. Harras and K. Almeroth, "[Scheduling Messengers in Disconnected Clustered Mobile Networks](#)," Ad Hoc & Sensor Wireless Networks, vol. 9, num. 3-4, pp. 275-304, March-April 2010.
46. A. Jardosh, K. Papagiannaki, E. Belding, K. Almeroth, G. Iannaccone, and B. Vinnakota, "[Green WLANs: On-Demand WLAN Infrastructures](#)," ACM Journal on Mobile Networks and Applications (MONET), vol. 14, num. 6, pp. 798-814, December 2009.
45. M. Wittie, K. Harras, K. Almeroth, and E. Belding, "[On the Implications of Routing Metric Staleness in Delay Tolerant Networks](#)," Computer Communications Special Issue on Delay and Disruption Tolerant Networking, vol. 32, num. 16, pp. 1699-1709, October 2009.
44. K. Harras, L. Deek, C. Holman, and K. Almeroth, "[DBS-IC: An Adaptive Data Bundling System for Intermittent Connectivity](#)," Computer Communications Special Issue on Delay and Disruption Tolerant Networking, vol. 32, num. 16, pp. 1687-1698, October 2009.
43. S. Karpinski, E. Belding, K. Almeroth, and J. Gilbert, "[Linear Representations of Network Traffic](#)," ACM Journal on Mobile Networks and Applications (MONET), vol. 14, num. 4, pp. 368-386, August 2009.
42. K. Harras and K. Almeroth, "[Controlled Flooding in Disconnected Sparse Mobile Networks](#)," Wireless Communications and Mobile Computing (WCMC) Journal, vol. 9, num. 1, pp. 21-33, January 2009.
41. R. Mayer, A. Stull, K. DeLeeuw, K. Almeroth, B. Bimber, D. Chun, M. Bulger, J. Campbell, A. Knight, and H. Zhang, "[Clickers in College Classrooms: Fostering Learning with Questioning Methods in Large Lecture Classes](#)," Contemporary Educational Psychology, vol. 34, num. 1, pp. 51-57, January 2009.
40. A. Knight, K. Almeroth, and B. Bimber, "[Design, Implementation and Deployment of PAIRwise](#)," Journal of Interactive Learning Research (JILR), vol. 19, num. 3, pp. 489-508, July 2008.
39. A. Garyfalos and K. Almeroth, "[Coupons: A Multilevel Incentive Scheme for Information Dissemination in Mobile Networks](#)," IEEE Transactions on Mobile Computing, vol. 7, num. 6, pp. 792-804, June 2008.
38. I. Sheriff, K. Ramachandran, E. Belding, and K. Almeroth, "[A Multi-Radio 802.11 Mesh Network Architecture](#)," ACM Journal on Mobile Networks and Applications (MONET), vol. 13, num. 1-2, pp. 132-146, April 2008.
37. M. Bulger, R. Mayer, K. Almeroth, and S. Blau, "[Measuring Learner Engagement in Computer-Equipped College Classrooms](#)," Journal of Educational Multimedia and Hypermedia, vol. 17, num. 2, pp. 129-143, April 2008.
36. G. Swamynathan, B. Zhao, and K. Almeroth, "[Exploring the Feasibility of Proactive Reputations](#)," Concurrency and Computation: Practice and Experience, vol. 20, num. 2, pp. 155-166, February 2008.
35. G. Swamynathan, B. Zhao, K. Almeroth, and H. Zheng, "[Globally Decoupled Reputations for Large Distributed Networks](#)," Advances in Multimedia, vol. 2007, pp. 1-14, 2007.

34. R. Mayer, A. Stull, J. Campbell, K. Almeroth, B. Bimber, D. Chun and A. Knight, "[Overestimation Bias in Self-reported SAT Scores](#)," *Educational Psychology Review*, vol. 19, num. 4, pp. 443-454, December 2007.

33. P. Namburi, K. Sarac and K. Almeroth, "[Practical Utilities for Monitoring Multicast Service Availability](#)," *Computer Communications Special Issue on Monitoring and Measurement of IP Networks*, vol. 29, num. 10, pp. 1675-1686, June 2006.

32. R. Chalmers, G. Krishnamurthi and K. Almeroth, "[Enabling Intelligent Handovers in Heterogeneous Wireless Networks](#)," *ACM Journal on Mobile Networks and Applications (MONET)*, vol. 11, num. 2, pp. 215-227, April 2006.

31. H. Lundgren, K. Ramachandran, E. Belding, K. Almeroth, M. Benny, A. Hewatt, A. Touma and A. Jardosh, "[Experience from the Design, Deployment and Usage of the UCSB MeshNet Testbed](#)," *IEEE Wireless Communications*, vol. 13, num. 2, pp. 18-29, April 2006.

30. R. Mayer, K. Almeroth, B. Bimber, D. Chun, A. Knight and A. Campbell, "[Technology Comes to College: Understanding the Cognitive Consequences of Infusing Technology in College Classrooms](#)," *Educational Technology*, vol. 46, num. 2, pp. 48-53, March-April 2006.

29. A. Garyfalos and K. Almeroth, "[A Flexible Overlay Architecture for Mobile IPv6 Multicast](#)," *Journal on Selected Areas in Communications (JSAC) Special Issue on Wireless Overlay Networks Based on Mobile IPv6*, vol. 23, num. 11, pp. 2194-2205, November 2005.

28. K. Sarac and K. Almeroth, "[Monitoring IP Multicast in the Internet: Recent Advances and Ongoing Challenges](#)," *IEEE Communications*, vol. 43, num. 10, pp. 85-91, October 2005.

27. K. Sarac and K. Almeroth, "[Application Layer Reachability Monitoring for IP Multicast](#)," *Computer Networks*, vol. 48, num. 2, pp. 195-213, June 2005.

26. A. Jardosh, E. Belding, K. Almeroth and S. Suri, "[Real-world Environment Models for Mobile Network Evaluation](#)," *Journal on Selected Areas in Communications Special Issue on Wireless Ad hoc Networks*, vol. 23, num. 3, pp. 622-632, March 2005.

25. S. Rollins and K. Almeroth, "[Evaluating Performance Tradeoffs in a One-to-Many Peer Content Distribution Architecture](#)," *Journal of Internet Technology*, vol. 5, num. 4, pp. 373-387, Fall 2004.

24. K. Sarac and K. Almeroth, "[Tracetree: A Scalable Mechanism to Discover Multicast Tree Topologies in the Network](#)," *IEEE/ACM Transactions on Networking*, vol. 12, num. 5, pp. 795-808, October 2004.

23. K. Sarac and K. Almeroth, "[A Distributed Approach for Monitoring Multicast Service Availability](#)," *Journal of Network and Systems Management*, vol. 12, num. 3, pp. 327-348, September 2004.

22. P. Rajvaidya, K. Ramachandran and K. Almeroth, "[Managing and Securing the Global Multicast Infrastructure](#)," *Journal of Network and Systems Management*, vol. 12, num. 3, pp. 297-326, September 2004.

21. P. Rajvaidya and K. Almeroth, "[Multicast Routing Instabilities](#)," *IEEE Internet Computing*, vol. 8, num. 5, pp. 42-49, September/October 2004.

20. D. Johnson, R. Patton, B. Bimber, K. Almeroth and G. Michaels, "[Technology and Plagiarism in the University: Brief Report of a Trial in Detecting Cheating](#)," *Association for the Advancement of Computing in Education (AACE) Journal*, vol. 12, num. 3, pp. 281-299, Summer 2004.

19. R. Chalmers and K. Almeroth, "[A Security Architecture for Mobility-Related Services](#)," *Journal of Wireless Personal Communications*, vol 29, num. 3, pp. 247-261, June 2004.
18. B. Stiller, K. Almeroth, J. Altmann, L. McKnight, and M. Ott, "[Pricing for Content in the Internet](#)," *Computer Communications*, vol. 27, num. 6, pp. 522-528, April 2004.
17. S. Rollins and K. Almeroth, "[Lessons Learned Deploying a Digital Classroom](#)," *Journal of Interactive Learning Research (JILR)*, vol. 15, num. 2, pp. 169-185, April 2004.
16. S. Jagannathan and K. Almeroth, "[A Dynamic Pricing Scheme for E-Content at Multiple Levels-of-Service](#)," *Computer Communications*, vol. 27, num. 4, pp. 374-385, March 2004.
15. K. Almeroth, "[Using Satellite Links in the Delivery of Terrestrial Multicast Traffic](#)," *Internetworking and Computing over Satellites*, Kluwer Academic Publishers, 2003.
14. R. Chalmers and K. Almeroth, "[On the Topology of Multicast Trees](#)," *IEEE/ACM Transactions on Networking*, vol. 11, num. 1, pp. 153-165, January 2003.
13. S. Jagannathan, J. Nayak, K. Almeroth, and M. Hofmann, "[On Pricing Algorithms for Batched Content Delivery Systems](#)," *Electronic Commerce Research and Applications Journal*, vol. 1, num. 3-4, pp. 264-280, Fall 2002.
12. D. Makofske and K. Almeroth, "[Multicast Sockets: Practical Guide for Programmers](#)," *Morgan Kaufmann Publishers*, November 2002.
11. S. Jagannathan and K. Almeroth, "[Price Issues in Delivering E-Content On-Demand](#)," *ACM Sigecom Exchanges*, vol. 3, num. 2, pp. 18-27, May 2002.
10. D. Makofske and K. Almeroth, "[From Television to Internet Video-on-Demand: Techniques and Tools for VCR-Style Interactivity](#)," *Software: Practice and Experience*, vol. 31, num. 8, pp. 781-801, July 2001.
9. K. Sarac and K. Almeroth, "[Supporting Multicast Deployment Efforts: A Survey of Tools for Multicast Monitoring](#)," *Journal on High Speed Networking*, Special Issue on Management of Multimedia Networking, vol. 9, num. 3/4, pp. 191-211, March 2001.
8. K. Almeroth, "[Adaptive, Workload-Dependent Scheduling for Large-Scale Content Delivery Systems](#)," *Transactions on Circuits and Systems for Video Technology, Special Issue on Streaming Video*, vol. 11, num. 3, pp. 426-439, March 2001.
7. D. Makofske and K. Almeroth, "[Real-Time Multicast Tree Visualization and Monitoring](#)," *Software: Practice and Experience*, vol. 30, num. 9, pp. 1047-1065, July 2000.
6. M. Ammar, K. Almeroth, R. Clark and Z. Fei, "Multicast Delivery of WWW Pages," *Electronic Commerce Technology Trends: Challenges and Opportunities*, IBM Press, February 2000.
5. K. Almeroth, "[The Evolution of Multicast: From the MBone to Inter-Domain Multicast to Internet Deployment](#)," *IEEE Network Special Issue on Multicasting*, vol. 10, num. 1, pp. 10-20, January/February 2000.
4. K. Almeroth and M. Ammar, "[An Alternative Paradigm for Scalable On-Demand Applications: Evaluating and Deploying the Interactive Multimedia Jukebox](#)," *IEEE Transactions on Knowledge and Data Engineering Special Issue on Web Technologies*, vol. 11, num. 4, pp 658-672, July/August 1999.

3. K. Almeroth and M. Ammar, "[The Interactive Multimedia Jukebox \(IMJ\): A New Paradigm for the On-Demand Delivery of Audio/Video](#)," *Computer Networks and ISDN Systems*, vol. 30, no. 1, April 1998.
2. K. Almeroth and M. Ammar, "[Multicast Group Behavior in the Internet's Multicast Backbone \(MBone\)](#)," *IEEE Communications*, vol. 35, no. 6, pp. 124-129, June 1997.
1. K. Almeroth and M. Ammar, "[On the Use of Multicast Delivery to Provide a Scalable and Interactive Video-on-Demand Service](#)," *Journal on Selected Areas of Communication (JSAC)*, vol. 14, no. 6, pp. 1110-1122, August 1996.

B. Conference Papers with Proceedings (refereed)

89. D. Havey and K. Almeroth, "[Active Sense Queue Management \(ASQM\)](#)," *IFIP Networking Conference*, Toulouse, FRANCE, May 2015.
88. L. Deek, E. Garcia-Villegas, E. Belding, S.J. Lee, and K. Almeroth, "[Joint Rate and Channel Width Adaptation in 802.11 MIMO Wireless Networks](#)," *IEEE Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON)*, New Orleans, LA, USA, June 2013.
87. D. Havey and K. Almeroth, "[Fast Wireless Protocol: A Network Stack Design for Wireless Transmission](#)," *IFIP Networking Conference*, Brooklyn, New York, USA, May 2013.
86. M. Tavakolifard, J. Gulla, K. Almeroth, J. Ingvaldsen, G. Nygreen, and E. Berg, "[Tailored News in the Palm of Your HAND: A Multi-Perspective Transparent Approach to News Recommendation](#)," *Demo Track at the International World Wide Web Conference (WWW)*, Rio de Janeiro, BRAZIL, May 2013.
85. S. Patterson, M. Wittie, K. Almeroth, and B. Bamieh, "[Network Optimization with Dynamic Demands and Link Prices](#)," *Allerton Conference*, Monticello, Illinois, USA, October 2012.
84. D. Havey, R. Chertov, and K. Almeroth, "[Receiver Driven Rate Adaptation](#)," *ACM Multimedia Systems Conference (MMSys)*, Chapel Hill, North Carolina, USA, February 2012.
83. M. Tavakolifard and K. Almeroth, "[Trust 2.0: Who to Believe in the Flood of Online Data?](#)," *International Conference on Computing, Networking and Communications (ICNC)*, Maui, Hawaii, USA, January 2012.
82. L. Deek, E. Garcia-Villegas, E. Belding, S.J. Lee, and K. Almeroth, "[The Impact of Channel Bonding on 802.11n Network Management](#)," *ACM CoNEXT*, Tokyo, JAPAN, December 2011.
81. L. Deek, X. Zhou, K. Almeroth, and H. Zheng, "[To Preempt or Not: Tackling Bid and Time-based Cheating in Online Spectrum Auctions](#)," *IEEE Infocom*, Shanghai, CHINA, April 2011.
80. M. Wittie, V. Pejovic, L. Deek, K. Almeroth, and B. Zhao, "[Exploiting Locality of Interest in Online Social Networks](#)," *ACM CoNEXT*, Philadelphia, Pennsylvania, USA, November 2010.
79. R. Chertov and K. Almeroth, "[Using BGP in a Satellite-Based Challenged Network Environment](#)," *IEEE Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON)*, Boston, Massachusetts, USA, June 2010.
78. R. Chertov, D. Havey and K. Almeroth, "[MSET: A Mobility Satellite Emulation Testbed](#)," *IEEE*

Infocom, San Diego, California, USA, March 2010.

77. B. Stone-Gross, A. Moser, C. Kruegel, E. Kirda, and K. Almeroth, "[FIRE: FInder Rogue nEtworks](#)," *Annual Computer Security Applications Conference (ACSAC)*, Honolulu, Hawaii, USA, December 2009.
76. M. Wittie, K. Almeroth, E. Belding, I. Rimac, and V. Hilt, "[Internet Service in Developing Regions Through Network Coding](#)," *IEEE Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON)*, Rome, ITALY, June 2009.
75. R. Chertov and K. Almeroth, "[High-Fidelity Link Shaping](#)," *International Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities (TRIDENTCOM)*, Washington DC, USA, April 2009.
74. L. Deek, K. Almeroth, M. Wittie, and K. Harras, "[Exploiting Parallel Networks Using Dynamic Channel Scheduling](#)," *International Wireless Internet Conference (WICON)*, Maui, Hawaii, USA, November 2008.
73. D. Havey, E. Barlas, R. Chertov, K. Almeroth, and E. Belding, "[A Satellite Mobility Model for QUALNET Network Simulations](#)," *IEEE Military Communications Conference (MILCOM)*, San Diego, California, USA, November 2008.
72. J. Kayfetz and K. Almeroth, "[Creating Innovative Writing Instruction for Computer Science Graduate Students](#)," *ASEE/IEEE Frontiers in Education (FIE) Conference*, Saratoga Springs, New York, USA, October 2008.
71. G. Swamynathan, B. Zhao, K. Almeroth, and S. Rao, "[Towards Reliable Reputations for Dynamic Networked Systems](#)," *IEEE International Symposium on Reliable Distributed Systems (SRDS)*, Napoli, ITALY, October 2008.
70. B. Stone-Gross, D. Sigal, R. Cohn, J. Morse, K. Almeroth, and C. Krugel, "[VeriKey: A Dynamic Certificate Verification System for Public Key Exchanges](#)," *Conference on Detection of Intrusions and Malware & Vulnerability Assessment (DIMVA)*, Paris, FRANCE, July 2008.
69. P. Acharya, A. Sharma, E. Belding, K. Almeroth, K. Papagiannaki, "[Congestion-Aware Rate Adaptation in Wireless Networks: A Measurement-Driven Approach](#)," *IEEE Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON)*, San Francisco, California, USA, June 2008.
68. A. Jardosh, P. Suwannat, T. Hollerer, E. Belding, and K. Almeroth, "[SCUBA: Focus and Context for Real-time Mesh Network Health Diagnosis](#)," *Passive and Active Measurement Conference (PAM)*, Cleveland, Ohio, USA, April 2008.
67. B. Stone-Gross, C. Wilson, K. Almeroth, E. Belding, H. Zheng, K. Papagiannaki, "[Malware in IEEE 802.11 Wireless Networks](#)," *Passive and Active Measurement Conference (PAM)*, Cleveland, Ohio, USA, April 2008.
66. R. Raghavendra, E. Belding, K. Papagiannaki, and K. Almeroth, "[Understanding Handoffs in Large IEEE 802.11 Wireless Networks](#)," *Internet Measurement Conference (IMC)*, San Diego, California, USA, October 2007.
65. M. Wittie, B. Stone-Gross, K. Almeroth and E. Belding, "[MIST: Cellular Data Network Measurement for Mobile Applications](#)," *IEEE International Conference on Broadband Communications, Networks,*

and Systems (BroadNets), Raleigh, North Carolina, USA, September 2007.

64. S. Karpinski, E. Belding, K. Almeroth, "[Wireless Traffic: The Failure of CBR Modeling](#)," *IEEE International Conference on Broadband Communications, Networks, and Systems (BroadNets)*, Raleigh, North Carolina, USA, September 2007.
63. A. Knight, K. Almeroth, H. Zhang, R. Mayer, and K. DeLeeuw, "[Data Cafe: A Dining Car Approach to Educational Research Data Management and Distribution](#)," *World Conference on Educational Multimedia, Hypermedia & Telecommunications (ED MEDIA)*, Vancouver, CANADA, June 2007.
62. H. Zhang, K. Almeroth, A. Knight, M. Bulger, and R. Mayer, "[Moodog: Tracking Students' Online Learning Activities](#)," *World Conference on Educational Multimedia, Hypermedia & Telecommunications (ED MEDIA)*, Vancouver, CANADA, June 2007.
61. M. Bulger, K. Almeroth, R. Mayer, D. Chun, A. Knight, H. Collins, "[Effects of Instructor Engagement on Student Use of a Course Management System](#)," Association for Psychological Science (APS) Annual Conference, Washington DC, USA, May 2007.
60. R. Mayer, A. Stull, K. Almeroth, B. Bimber, D. Chun, M. Bulger, J. Campbell, Allan Knight, and H. Zhang, "[Using Technology-Based Methods to Foster Learning in Large Lecture Classes: Evidence for the Pedagogic Value of Clickers](#)," *American Educational Research Association (AERA) Annual Conference*, Chicago, Illinois, USA, April 2007.
59. K. Ramachandran, I. Sheriff, E. Belding, and K. Almeroth, "[Routing Stability in Static Wireless Mesh Networks](#)," *Passive and Active Measurement Conference (PAM)*, Louvain-la-Neuve, BELGIUM, April 2007.
58. G. Swamynathan, T. Close, S. Banerjee, R. McGeer, B. Zhao, and K. Almeroth, "[Scalable Access Control For Web Services](#)," *International Conference on Creating, Connecting and Collaborating through Computing (C5)*, Kyoto, JAPAN, January 2007.
57. A. Knight, M. Bulger, K. Almeroth, and H. Zhang, "[Is Learning Really a Phone Call Away? Knowledge Transfer in Mobile Learning](#)," *World Conference on Mobile Learning (mLearn)*, Banff, Alberta, CANADA, October 2006.
56. J. Kurian, K. Sarac, and K. Almeroth, "[Defending Network-Based Services Against Denial of Service Attacks](#)," *International Conference on Computer Communication and Networks (IC3N)*, Arlington, Virginia, USA, October 2006.
55. A. Jardosh, K. Sanzgiri, E. Belding and K. Almeroth, "[IQU: Practical Queue-Based User Association Management for WLANs--Case Studies, Architecture, and Implementation](#)," ACM Mobicom, Marina del Rey, California, USA, September 2006.
54. C. Holman, K. Harras, and K. Almeroth, "[A Proactive Data Bundling System for Intermittent Mobile Connections](#)," *IEEE International Conference on Sensor and Ad Hoc Communications and Networks (SECON)*, Reston, Virginia, USA, September 2006.
53. G. Banks, M. Cova, V. Felmetzger, K. Almeroth, R. Kemmerer and G. Vigna, "[SNOOZE: toward a Stateful NetwOrk prOtocol fuzZEr](#)," *Information Security Conference (ISC)*, Samos Island, GREECE, September 2006.
52. K. Harras and K. Almeroth, "[Inter-Regional Messenger Scheduling in Delay Tolerant Mobile Networks](#)," *IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks*

(WoWMoM), Niagara Falls, New York, USA, June 2006.

51. M. Bulger, R. Mayer, and K. Almeroth, "[Engaged By Design: Using Simulation to Promote Active Learning](#)," **Outstanding Paper** at the *World Conference on Educational Multimedia, Hypermedia & Telecommunications (ED MEDIA)*, Orlando, Florida, USA, June 2006.
50. A. Knight, K. Almeroth, R. Mayer, D. Chun, and B. Bimber, "[Observations and Recommendations for Using Technology to Extend Interaction](#)," *World Conference on Educational Multimedia, Hypermedia & Telecommunications (ED MEDIA)*, Orlando, Florida, USA, June 2006.
49. H. Zhang, and K. Almeroth, "[A Simple Classroom Network Access Control System](#)," *World Conference on Educational Multimedia, Hypermedia & Telecommunications (ED MEDIA)*, Orlando, Florida, USA, June 2006.
48. K. Harras and K. Almeroth, "[Transport Layer Issues in Delay Tolerant Mobile Networks](#)," *IFIP Networking Conference*, Coimbra, PORTUGAL, May 2006.
47. R. Mayer, A. Stull, J. Campbell, K. Almeroth, B. Bimber, D. Chun and A. Knight, "[Some Shortcomings of Soliciting Students' Self-Reported SAT Scores](#)," *American Educational Research Association (AERA) Annual Conference*, San Francisco, California, USA, April 2006.
46. K. Ramachandran, E. Belding, K. Almeroth, and M. Buddhikot, "[Interference-Aware Channel Assignment in Multi-Radio Wireless Mesh Networks](#)," *IEEE Infocom*, Barcelona, SPAIN, April 2006.
45. A. Jardosh, K. Ramachandran, K. Almeroth, and E. Belding, "[Understanding Congestion in IEEE 802.11b Wireless Networks](#)," *ACM/USENIX Internet Measurement Conference (IMC)*, Berkeley, California, USA, October 2005.
44. H. Zhang, K. Almeroth and M. Bulger, "[An Activity Monitoring System to Support Classroom Research](#)," *World Conference on Educational Multimedia, Hypermedia & Telecommunications (ED MEDIA)*, Montreal, Quebec, CANADA, pp. 1444-1449, June 2005.
43. Z. Xiang, H. Zhang, J. Huang, S. Song and K. Almeroth, "[A Hidden Environment Model for Constructing Indoor Radio Maps](#)," *IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, Taormina, ITALY, June 2005.
42. K. Harras, K. Almeroth and E. Belding, "[Delay Tolerant Mobile Networks \(DTMNs\): Controlled Flooding in Sparse Mobile Networks](#)," *IFIP Networking Conference*, Waterloo, Ontario, CANADA, May 2005.
41. A. Garyfalos and K. Almeroth, "[Coupons: Wide Scale Information Distribution for Wireless Ad Hoc Networks](#)," *IEEE Global Telecommunications Conference (Globecom) Global Internet and Next Generation Networks Symposium*, Dallas, Texas, USA, pp. 1655-1659, December 2004.
40. A. Knight and K. Almeroth, "[DeCAF: A Digital Classroom Application Framework](#)," *IASTED International Conference on Communications, Internet and Information Technology (CIIT)*, St. Thomas, US Virgin Islands, November 2004.
39. P. Namburi, K. Sarac and K. Almeroth, "[SSM-Ping: A Ping Utility for Source Specific Multicast](#)," *IASTED International Conference on Communications, Internet and Information Technology (CIIT)*, St. Thomas, US Virgin Islands, November 2004.
38. K. Ramachandran, E. Belding and K. Almeroth, "[DAMON: A Distributed Architecture for Monitoring](#)

[Multi-hop Mobile Networks](#)," *IEEE International Conference on Sensor and Ad Hoc Communications and Networks (SECON)*, Santa Clara, California, USA, October 2004.

37. A. Garyfalos and K. Almeroth, "[Coupon Based Incentive Systems and the Implications of Equilibrium Theory](#)," *IEEE Conference on Electronic Commerce (CEC)*, San Diego, California, USA, pp. 213-220, July 2004.
36. A. Knight, K. Almeroth and B. Bimber, "[An Automated System for Plagiarism Detection Using the Internet](#)," *World Conference on Educational Multimedia, Hypermedia & Telecommunications (ED MEDIA)*, Lugano, Switzerland, pp. 3619-3625, June 2004.
35. H. Zhang and K. Almeroth, "[Supplement to Distance Learning: Design for a Remote TA Support System](#)," *World Conference on Educational Multimedia, Hypermedia & Telecommunications (ED MEDIA)*, Lugano, Switzerland, pp. 2821-2830, June 2004.
34. U. Mohan, K. Almeroth and E. Belding, "[Scalable Service Discovery in Mobile Ad hoc Networks](#)," *IFIP Networking Conference*, Athens, Greece, pp. 137-149, May 2004.
33. V. Thanedar, K. Almeroth and E. Belding, "[A Lightweight Content Replication Scheme for Mobile Ad hoc Environments](#)," *IFIP Networking Conference*, Athens, Greece, pp. 125-136, May 2004.
32. R. Chalmers and K. Almeroth, "[A Mobility Gateway for Small-Device Networks](#)," *IEEE International Conference on Pervasive Computing and Communications (PerCom)*, Orlando, Florida, USA, March 2004.
31. A. Jardosh, E. Belding, K. Almeroth and S. Suri, "[Towards Realistic Mobility Models For Mobile Ad hoc Networks](#)," *ACM Mobicom*, San Diego, California, USA, September 2003.
30. K. Sarac, P. Namburi and K. Almeroth, "[SSM Extensions: Network Layer Support for Multiple Senders in SSM](#)," *International Conference on Computer Communication and Networks (IC3N)*, Dallas, Texas, USA, October 2003.
29. K. Ramachandran and K. Almeroth, "[MAFIA: A Multicast Management Solution for Access Control and Traffic Filtering](#)," *IEEE/IFIP Conference on Management of Multimedia Networks and Services (MMNS)*, Belfast, Northern Ireland, September 2003.
28. J. Humfrey, S. Rollins, K. Almeroth, and B. Bimber, "[Managing Complexity in a Networked Learning Environment](#)," *World Conference on Educational Multimedia, Hypermedia & Telecommunications (ED MEDIA)*, Honolulu, Hawaii, USA, pp. 60-63, June 2003.
27. K. Almeroth, S. Rollins, Z. Shen, and B. Bimber, "[Creating a Demarcation Point Between Content Production and Encoding in a Digital Classroom](#)," *World Conference on Educational Multimedia, Hypermedia & Telecommunications (ED MEDIA)*, Honolulu, Hawaii, USA, pp. 2-5, June 2003.
26. M. Kolsch, K. Kvilekval, and K. Almeroth, "[Improving Speaker Training with Interactive Lectures](#)," *World Conference on Educational Multimedia, Hypermedia & Telecommunications (ED MEDIA)*, Honolulu, Hawaii, USA, June 2003.
25. P. Rajvaidya and K. Almeroth, "[Analysis of Routing Characteristics in the Multicast Infrastructure](#)," *IEEE Infocom*, San Francisco, California, USA, April 2003.
24. S. Rollins and K. Almeroth, "[Pixie: A Jukebox Architecture to Support Efficient Peer Content Exchange](#)," *ACM Multimedia*, Juan Les Pins, FRANCE, December 2002.

23. S. Rollins, R. Chalmers, J. Blanquer, and K. Almeroth, "[The Active Information System\(AIS\): A Model for Developing Scalable Web Services](#)," *IASTED International Conference on Internet and Multimedia Systems and Applications (IMSA)*, Kauai, Hawaii, USA, August 2002.

22. S. Rollins and K. Almeroth, "[Seminal: Additive Semantic Content for Multimedia Streams](#)," *IASTED International Conference on Internet and Multimedia Systems and Applications (IMSA)*, Kauai, Hawaii, USA, August 2002.

21. B. Stiller, K. Almeroth, J. Altmann, L. McKnight, and M. Ott, "[Content Pricing in the Internet](#)," *SPIE ITCom Conference on Internet Performance and Control of Network Systems (IPCNS)*, Boston, Massachusetts, USA, July 2002.

20. S. Jagannathan, J. Nayek, K. Almeroth and M. Hofmann, "[A Model for Discovering Customer Value for E-Content](#)," *ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD)*, Edmonton, Alberta, CANADA, July 2002.

19. S. Rollins and K. Almeroth, "[Deploying and Infrastructure for Technologically Enhanced Learning](#)," **Outstanding Paper** at the *World Conference on Educational Multimedia, Hypermedia & Telecommunications (ED MEDIA)*, Denver, Colorado, USA, pp. 1651-1656, June 2002.

18. P. Rajvaidya and K. Almeroth, "[Building the Case for Distributed Global Multicast Monitoring](#)," *Multimedia Computing and Networking (MMCN)*, San Jose, California, USA, January 2002.

17. S. Jagannathan and K. Almeroth, "[An Adaptive Pricing Scheme for Content Delivery Systems](#)," *IEEE Global Internet*, San Antonio, Texas, USA, November 2001.

16. K. Sarac and K. Almeroth, "[Providing Scalable Many-to-One Feedback in Multicast Reachability Monitoring Systems](#)," *IFIP/IEEE International Conference on Management of Multimedia Networks and Services (MMNS)*, Chicago, Illinois, USA, October 2001.

15. S. Jagannathan and K. Almeroth, "[The Dynamics of Price, Revenue and System Utilization](#)," *IFIP/IEEE International Conference on Management of Multimedia Networks and Services (MMNS)*, Chicago, Illinois, USA, October 2001.

14. A. Kanwar, K. Almeroth, S. Bhattacharyya, and M. Davy, "[Enabling End-User Network Monitoring via the Multicast Consolidated Proxy Monitor](#)," *SPIE ITCom Conference on Scalability and Traffic Control in IP Networks (STCIPN)*, Denver, Colorado, USA, August 2001.

13. S. Jagannathan and K. Almeroth, "[Using Tree Topology for Multicast Congestion Control](#)," *International Conference on Parallel Processing (ICPP)*, Valencia, SPAIN, September 2001.

12. P. Rajvaidya and K. Almeroth, "[A Router-Based Technique for Monitoring the Next-Generation of Internet Multicast Protocols](#)," *International Conference on Parallel Processing (ICPP)*, Valencia, SPAIN, September 2001.

11. R. Chalmers and K. Almeroth, "[Modeling the Branching Characteristics and Efficiency Gains of Global Multicast Trees](#)," *IEEE Infocom*, Anchorage, Alaska, USA, April 2001.

10. R. Chalmers and K. Almeroth, "[Developing a Multicast Metric](#)," *Global Internet*, San Francisco, California, USA, December 2000.

9. K. Sarac and K. Almeroth, "[Monitoring Reachability in the Global Multicast Infrastructure](#)," *IEEE International Conference on Network Protocols (ICNP)*, Osaka, JAPAN, November 2000.

8. K. Almeroth, "[A Long-Term Analysis of Growth and Usage Patterns in the Multicast Backbone \(MBone\)](#)," *IEEE INFOCOM*, Tel Aviv, ISRAEL, March 2000.
7. K. Almeroth, K. Obraczka and D. De Lucia, "[A Lightweight Protocol for Interconnecting Heterogeneous Devices in Dynamic Environments](#)," *IEEE International Conference on Multimedia Computing and Systems (ICMCS)*, Florence, ITALY, June 1999.
6. K. Almeroth and M. Ammar, "[The Interactive Multimedia Jukebox \(IMJ\): A New Paradigm for the On-Demand Delivery of Audio/Video](#)," **Best Paper** at the *Seventh International World Wide Web Conference (WWW)*, Brisbane, AUSTRALIA, April 1998.
5. K. Almeroth, M. Ammar and Z. Fei, "[Scalable Delivery of Web Pages Using Cyclic Best-Effort \(UDP\) Multicast](#)," *IEEE INFOCOM*, San Francisco, California, USA, June 1998.
4. K. Almeroth and M. Ammar, "[Delivering Popular Web Pages Using Cyclic Unreliable Multicast \(Extended Abstract\)](#)," *SPIE Conference on Voice, Video and Data Communications*, Dallas, Texas, USA, November 1997.
3. K. Almeroth, A. Dan, D. Sitaram and W. Tetzlaff, "[Long Term Resource Allocation in Video Delivery Systems](#)," *IEEE INFOCOM*, Kobe, JAPAN, April 1997.
2. K. Almeroth and M. Ammar, "[On the Performance of a Multicast Delivery Video-On-Demand Service with Discontinuous VCR Actions](#)," *International Conference on Communications (ICC)*, Seattle, Washington, USA, June 1995.
1. K. Almeroth and M. Ammar, "[A Scalable, Interactive Video-On-Demand Service Using Multicast Communication](#)," *International Conference on Computer Communication and Networks (IC3N)*, San Francisco, California, USA, September 1994.

C. Workshop Papers (refereed)

34. M. Tavakolifard, J. Gulla, K. Almeroth, F. Hopfgartner, B. Kille, T. Plumbaum, A. Lommatsch, T. Brodt, A. Bucko, and T. Heintz, "[Workshop and Challenge on News Recommender Systems](#)," *ACM RecSys News Recommender Systems (NRS) Workshop and Challenge*, Hong Kong, CHINA, October 2013.
33. M. Tavakolifard, K. Almeroth, and J. Gulla, "[Does Social Contact Matter? Modelling the Hidden Web of Trust Underlying Twitter](#)," *ACM International Workshop on Social Recommender Systems (SRS)*, Rio de Janeiro, BRAZIL, May 2013.
32. D. Johnson, E. Belding, K. Almeroth and G. van Stam, "[Internet Usage and Performance Analysis of a Rural Wireless Network in Macha, Zambia](#)," *ACM Networked Systems for Developing Regions (NSDR) Workshop*, San Francisco, California, USA, June 2010.
31. D. Havey, R. Chertov, and K. Almeroth, "[Wired Wireless Broadcast Emulation](#)," *International Workshop on Wireless Network Measurement (WiNMee)*, Seoul, Korea, June 2009.
30. R. Raghavendra, P. Acharya, E. Belding, and K. Almeroth, "[MeshMon: A Multi-Tiered Framework for Wireless Mesh Network Monitoring](#)," *ACM MobiHoc Wireless of the Students, by the Students, for the Students Workshop (S3)*, New Orleans, Louisiana, USA, May 2009.

29. G. Swamynathan, C. Wilson, B. Boe, B. Zhao, and K. Almeroth, "[Do Social Networks Improve e-Commerce: A Study on Social Marketplaces](#)," *ACM Sigcomm Workshop on Online Social Networks (WOSN)*, Seattle, Washington, USA, August 2008.

28. R. Raghavendra, E. Belding, and K. Almeroth, "[Antler: A Multi-Tiered Approach to Automated Wireless Network Management](#)," *IEEE Workshop on Automated Network Management (ANM)*, Phoenix, Arizona, USA, April 2008.

27. S. Karpinski, E. Belding, and K. Almeroth, "[Towards Realistic Models of Wireless Workload](#)," *International Workshop on Wireless Network Measurement (WiNMee)*, Limassol, CYPRUS, April 2007.

26. K. Harras, M. Wittie, K. Almeroth, and E. Belding, "[ParaNets: A Parallel Network Architecture for Challenged Networks](#)," *IEEE Workshop on Mobile Computing Systems and Applications (HotMobile)*, Tucson, Arizona, USA, February 2007.

25. H. Caiuiro-Monge, K. Almeroth, M. del Mar Alvarez-Rohena, "[Friend Relay: A Resource Sharing Framework for Mobile Wireless Devices](#)," *ACM International Workshop on Wireless Mobile Applications and Services on WLAN Hotspots (WMASH)*, Los Angeles, California, September 2006.

24. G. Swamynathan, Ben Y. Zhao and K. Almeroth, "[Exploring the Feasibility of Proactive Reputations](#)," *International Workshop on Peer-to-Peer Systems (IPTPS)*, Santa Barbara, California, USA, February 2006.

23. G. Swamynathan, Ben Y. Zhao and K. Almeroth, "[Decoupling Service and Feedback Trust in a Peer-to-Peer Reputation System](#)," *International Workshop on Applications and Economics of Peer-to-Peer Systems (AEPP)*, Nanjing, CHINA, November 2005.

22. K. Ramachandran, M. Buddhikot, G. Chandranmenon, S. Miller, E. Belding, and K. Almeroth, "[On the Design and Implementation of Infrastructure Mesh Networks](#)," *IEEE Workshop on Wireless Mesh Networks (WiMesh)*, Santa Clara, California, USA, September 2005.

21. A. Jardosh, K. Ramachandran, K. Almeroth and E. Belding, "[Understanding Link-Layer Behavior in Highly Congested IEEE 802.11b Wireless Networks](#)," *Sigcomm Workshop on Experimental Approaches to Wireless Network Design and Analysis (EWIND)*, Philadelphia, Pennsylvania, USA, August 2005.

20. A. Sen Mazumder, K. Almeroth and K. Sarac, "[Facilitating Robust Multicast Group Management](#)," *Network and Operating System Support for Digital Audio and Video (NOSSDAV)*, Skamania, Washington, USA, June 2005.

19. Y. Sun, I. Sheriff, E. Belding and K. Almeroth, "[An Experimental Study of Multimedia Traffic Performance in Mesh Networks](#)," *MobiSys International Workshop on Wireless Traffic Measurements and Modeling (WitMeMo)*, Seattle, Washington, USA, June 2005.

18. K. Ramachandran, K. Almeroth and E. Belding, "[A Framework for the Management of Large-Scale Wireless Network Testbeds](#)," *International Workshop on Wireless Network Measurement (WiNMee)*, Trentino, ITALY, April 2005.

17. A. Garyfalos, K. Almeroth and K. Sanzgiri, "[Deployment Complexity Versus Performance Efficiency in Mobile Multicast](#)," *International Workshop on Broadband Wireless Multimedia: Algorithms, Architectures and Applications (BroadWiM)*, San Jose, California, USA, October 2004.

16. C. Ho, K. Ramachandran, K. Almeroth and E. Belding, "[A Scalable Framework for Wireless Network Monitoring](#)," *ACM International Workshop on Wireless Mobile Applications and Services on WLAN Hotspots (WMASH)*, Philadelphia, Pennsylvania, USA, October 2004.
15. A. Garyfalos, K. Almeroth and J. Finney, "[A Hybrid of Network and Application Layer Multicast for Mobile IPv6 Networks](#)," *International Workshop on Large-Scale Group Communication (LSGC)*, Florence, ITALY, October 2003.
14. A. Garyfalos, K. Almeroth and J. Finney, "[A Comparison of Network and Application Layer Multicast for Mobile IPv6 Networks](#)," *ACM Workshop on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*, San Diego, California, USA, September 2003.
13. S. Jagannathan, and K. Almeroth, "[Pricing and Resource Provisioning for Delivering E-Content On-Demand with Multiple Levels-of-Service](#)," *International Workshop on Internet Charging and QoS Technologies (ICQT)*, Zurich, SWITZERLAND, October 2002.
12. S. Rollins, K. Almeroth, D. Milojicic, and K. Nagaraja, "[Power-Aware Data Management for Small Devices](#)," *Workshop on Wireless Mobile Multimedia (WoWMoM)*, Atlanta, GA, USA, September 2002.
11. K. Almeroth, S. Bhattacharyya, and C. Diot, "[Challenges of Integrating ASM and SSM IP Multicast Protocol Architectures](#)," *International Workshop on Digital Communications: Evolutionary Trends of the Internet (IWDC)*, Taormina, ITALY, September 2001.
10. K. Sarac and K. Almeroth, "[Scalable Techniques for Discovering Multicast Tree Topology](#)," *Network and Operating System Support for Digital Audio and Video (NOSSDAV)* , Port Jefferson, New York, USA, June 2001.
9. P. Rajvaidya, K. Almeroth and K. Claffy, "[A Scalable Architecture for Monitoring and Visualizing Multicast Statistics](#)," *IFIP/IEEE International Workshop on Distributed Systems: Operations & Management (DSOM)*, Austin, Texas, USA, December 2000.
8. S. Jagannathan, K. Almeroth and A. Acharya, "[Topology Sensitive Congestion Control for Real-Time Multicast](#)," *Network and Operating System Support for Digital Audio and Video (NOSSDAV)* , Chapel Hill, North Carolina, USA, June 2000.
7. K. Sarac and K. Almeroth, "[Supporting the Need for Inter-Domain Multicast Reachability](#)," *Network and Operating System Support for Digital Audio and Video (NOSSDAV)* , Chapel Hill , North Carolina, USA, June 2000.
6. D. Makofske and K. Almeroth, "[MHealth: A Real-Time Multicast Tree Visualization and Monitoring Tool](#)," *Network and Operating System Support for Digital Audio and Video (NOSSDAV)* , Basking Ridge New Jersey, USA, June 1999.
5. K. Almeroth and Y. Zhang, "[Using Satellite Links as Delivery Paths in the Multicast Backbone \(MBone\)](#)," *ACM/IEEE International Workshop on Satellite-Based Information Services (WOSBIS)*, Dallas, Texas, USA, October 1998.
4. M. Ammar, K. Almeroth, R. Clark and Z. Fei, "[Multicast Delivery of Web Pages OR How to Make Web Servers Pushy](#)," *Workshop on Internet Server Performance (WISP)*, Madison, Wisconsin, USA, June 1998.
3. K. Almeroth and M. Ammar, "[Prototyping the Interactive Multimedia Jukebox](#)," *Mini-conference on Multimedia Appliances, Interfaces, and Trials--International Conference on Communications (ICC)*,

Montreal, Quebec, CANADA, June 1997.

2. K. Almeroth and M. Ammar, "[Collection and Modeling of the Join/Leave Behavior of Multicast Group Members in the MBone](#)," *High Performance Distributed Computing Focus Workshop (HPDC)*, Syracuse, New York, USA, August 1996.
1. K. Almeroth and M. Ammar, "[The Role of Multicast Communication in the Provision of Scalable and Interactive Video-On-Demand Service](#)," *Network and Operating System Support for Digital Audio and Video (NOSSDAV)* , Durham, New Hampshire, USA, April 1995.

D. Non-Refereed Publications

8. K. Almeroth, E. Belding, M. Buddhikot, G. Chandranmenon, S. Miller, and K. Ramachandran, "[Infrastructure Mesh Networks](#)," *U.S. Patent Application US20070070959 A1*, September 2005.
7. K. Almeroth, R. Caceres, A. Clark, R. Cole, N. Duffield, T. Friedman, K. Hedayat, K. Sarac, M. Westerlund, "[RTP Control Protocol Extended Reports \(RTCP XR\)](#)," *Internet Engineering Task Force (IETF) Request for Comments (RFC) 3611*, November 2003.
6. Z. Albanna, K. Almeroth, D. Meyer, and M. Schipper, "[IANA Guidelines for IPv4 Multicast Address Allocation](#)," *Internet Engineering Task Force (IETF) Request for Comments (RFC) 3171*, August 2001.
5. B. Quinn and K. Almeroth, "[IP Multicast Applications: Challenges and Solutions](#)," *Internet Engineering Task Force (IETF) Request for Comments (RFC) 3170*, September 2001.
4. K. Almeroth, L. Wei and D. Farinacci, "[Multicast Reachability Monitor \(MRM\) Protocol](#)," *Internet Engineering Task Force Internet Draft*, July 2000.
3. K. Almeroth and L. Wei, "[Justification for and use of the Multicast Reachability Monitor \(MRM\) Protocol](#)," *Internet Engineering Task Force Internet Draft*, March 1999.
2. K. Almeroth, "[Managing IP Multicast Traffic: A First Look at the Issues, Tools, and Challenges](#)," IP Multicast Initiative White Paper, San Jose, California, USA, February 1999.
1. K. Almeroth, K. Obraczka and D. De Lucia, "[Pseudo-IP: Providing a Thin Network Protocol for Semi-Intelligent Wireless Devices](#)," *DARPA/NIST Smart Spaces Workshop*, Gaithersburg, Maryland, USA, July 1998.

E. Released Software Systems

19. ***A Multi-radio Wireless Mesh Network Architecture -- <http://moment.cs.ucsbg.edu/tic/>***. Released December 1, 2006 (with K. Ramachandran, I. Sheriff, and E. Belding). The software as part of a multi-radio wireless mesh network that includes a Split Wireless Router that alleviates the interference that can occur between commodity radios within a single piece of hardware. The second is server software to perform channel assignment and communicate the assignments throughout the mesh network.
18. ***AODV-Spanning Tree (AODV-ST) -- <http://www.cs.ucsbg.edu/~krishna/aodv-st/>***. Released September 1, 2006 (with K. Ramachandran and E. Belding). AODV-ST is an extension of the well-known AODV

protocol specifically designed for wireless mesh networks. The advantages of AODV-ST over AODV include support for high throughput routing metrics, automatic route maintenance for common-case traffic, and low route discovery latency.

17. ***The Multicast Detective*** -- http://www.nmsl.cs.ucsbs.edu/mcast_detective/. Released September 1, 2005 (with A. Sen Mazumder). The multicast detective is a robust solution to determine the existence and nature of multicast service for a particular user. By performing a series of tests, a user can determine whether there is network support for multicast, and consequently, whether a multicast group join is likely to succeed.
16. ***AutoCap: Automatic and Accurate Captioning*** -- <http://www.nmsl.cs.ucsbs.edu/autocap/>. Released August 1, 2005 (with A. Knight). AutoCap is a software system that takes as input an audio/video file and a text transcript. AutoCap creates captions by aligning the utterances in the audio/video file to the transcript. For those words that are not recognized, AutoCap estimates when the words were spoken along with an error bound that gives the content creator an idea of caption accuracy. The result is a collection of accurately time-stamped captions that can be displayed with the video.
15. ***PAIRwise Plagiarism Detection System*** -- <http://cits.ucsbs.edu/pair/>. Released July 1, 2005 (with A. Knight). PAIRwise is a plagiarism detection system with: (1) an easy-to-use interface for submitting papers, (2) a flexible comparison engine that allows intra-class, inter-class, and Internet-based comparisons, and (3) an intuitive graphical presentation of results.
14. ***DAMON Multi-Hop Wireless Network Monitoring*** -- <http://moment.cs.ucsbs.edu/damon/>. Released October 1, 2004 (with K. Ramachandran and E. Belding). DAMON is a distributed system for monitoring multi-hop mobile networks. DAMON uses agents within the network to monitor network behavior and send collected measurements to data repositories. DAMON's generic architecture supports the monitoring of a wide range of protocol, device, or network parameters.
13. ***Multicast Firewall*** -- <http://www.nmsl.cs.ucsbs.edu/mafia/>. Released June 1, 2004 (with K. Ramachandran). MAFIA, a multicast firewall and traffic management solution, has the specific aim of strengthening multicast security through multicast access control, multicast traffic filtering, and DoS attack prevention.
12. ***AODV@IETF Peer Routing Software*** -- <http://moment.cs.ucsbs.edu/aodv-ietf/>. Released November 1, 2003 (with K. Ramachandran and E. Belding). One of the first large-scale efforts to run the Ad hoc On demand Distance Vector (AODV) routing protocol in a public space (at the Internet Engineering Task Force (IETF)). The implementation includes a daemon that runs on both the Linux and Windows operating systems.
11. ***Mobility Obstacles*** -- <http://moment.cs.ucsbs.edu/mobility/>. Released September 1, 2003 (with A. Jardosh, E. Belding, and S. Suri). The topology and movement of nodes in ad hoc protocol simulation are key factors in protocol performance. In this project, we have developed ns-2 simulation plug-ins that create more realistic movement models through the incorporation of obstacles. These obstacles are utilized to restrict both node movement and wireless transmissions.
10. ***mwalk*** -- <http://www.nmsl.cs.ucsbs.edu/mwalk/>. Released December 1, 2000 (with R. Chalmers). Mwalk is a collection of Java applications and Perl scripts which re-create a global view of a multicast session from mtrace and RTCP logs. Users to the site can download mwalk, examine the results of our analysis, or download data sets for use in simulations dependent on multicast tree characteristics.
9. ***MANTRA2*** -- <http://www.nmsl.cs.ucsbs.edu/mantra/>. Released December 1, 1999 (with P. Rajvaidya). This new version of MANTRA focuses on the visualization of inter-domain routing statistics. Working

in conjunction with the Cooperative Association for Internet Data Analysis (CAIDA) we have developed advanced collection and visualization techniques.

8. **MRM** -- <http://www.nmsl.cs.ucsbs.edu/mrm/>. Released October 1, 1999 (with K. Sarac). MRM is the Multicast Reachability Protocol. We have implemented an end-host agent that responds to MRM Manager commands. Our end-host agent works in conjunction with Cisco routers to detect and isolate multicast faults.
7. **MANTRA** -- <http://www.nmsl.cs.ucsbs.edu/mantra/>. Released January 1, 1999 (with P. Rajvaidya). MANTRA is the Monitoring and Analysis of Traffic in Multicast Routers. It uses scripts to collect and display data from backbone multicast routers.
6. **SDR Monitor** -- <http://www.nmsl.cs.ucsbs.edu/sdr-monitor/>. Released January 1, 1999 (with K. Sarac). The SDR Monitor receives e-mail updates from participants containing information about observed sessions in the MBone. A global view of multicast reachability is then constructed.
5. **The MHealth tool** -- <http://www.nmsl.cs.ucsbs.edu/mhealth/>. Released September 1, 1998 (with D. Makofske). The mhealth tool graphically visualizes MBone multicast group trees and provides 'health' information including end-to-end losses per receiver and losses on a per hop basis. The implementation required expertise in Java, the MBone tools, and Unix.
4. **The MControl tool** -- <http://www.nmsl.cs.ucsbs.edu/mcontrol/>. Released August 1, 1998 (with D. Makofske). Mcontrol is a tool to provide VCR-based interactivity for live MBone sessions. The implementation required expertise in Java, the MBone tools, and Unix.
3. **Interactive Multimedia Jukebox (IMJ)** -- <http://imj.ucsbs.edu/>. Released October 1, 1996. The IMJ combines the WWW and the MBone conferencing tools to provide a multi-channel video jukebox offering both instructional and entertainment programming on a wide scale. The implementation required expertise in HTML, Perl, C, the MBone tools, and Unix.
2. **Mlisten** -- <http://www.cc.gatech.edu/computing/Telecomm/mbone/>. Released September 1, 1995. A tool to continuously collect MBone multicast group membership information including number and location of members, membership duration, and inter-arrival time for all audio and video sessions. The implementation required expertise in C, Tcl/Tk, the MBone tools, and UNIX socket programming.
1. **Audio-on-Demand (AoD)**. March 1, 1995. A server/client prototype to demonstrate interactivity in near VoD systems. The AoD server provides songs-on-demand and VCR-like functions via multicast IP over Ethernet. The implementation required expertise in C, OpenWindows programming, UNIX socket programming, and network programming.

F. Tutorials, Panels and Invited Talks

- "25th Anniversary Panel," Network and Operating System Support for Digital Audio and Video (NOSSDAV), Portland, Oregon, USA, March 2015.
- "Sensing and Opportunistic Delivery of Ubiquitous Video in Health Monitoring, On-Campus and Social Network Applications," Workshop on Mobile Video Delivery (MoViD), Chapel Hill North Carolina, USA, February 2012.
- "Medium Access in New Contexts: Reinventing the Wheel?," USC Invited Workshop on Theory and

Practice in Wireless Networks, Los Angeles, California, USA, May 2008.

- "The Wild, Wild West: Wireless Networks Need a New Sheriff," University of Florida CISE Department Lecture Series, Gainesville, Florida, USA, February 2008.
- "Distinguishing Between Connectivity, Intermittent Connectivity, and Intermittent Disconnectivity," Keynote at the ACM MobiCom Workshop on Challenged Networks (CHANTS), Montreal, CANADA, September 2007.
- "The Three Ghosts of Multicast: Past, Present, and Future," Keynote at the Trans-European Research and Education Networking Association (TERENA) Networking Conference, Lynby, DENMARK, May 2007.
- "Multicast Help Wanted: From Where and How Much?," Keynote at the Workshop on Peer-to-Peer Multicasting (P2PM), Las Vegas, Nevada, USA, January 2007.
- "The Confluence of Wi-Fi and Apps: What to Expect Next," Engineering Insights, UC-Santa Barbara, Santa Barbara, California, USA, October 2006.
- "Challenges, Opportunities, and Implications for the Future Internet," University of Minnesota Digital Technology Center, Minneapolis, Minnesota, USA, September 2006.
- "Wireless Technology as a Catalyst: Possibilities for Next-Generation Interaction," Santa Barbara Forum on Digital Transitions, Santa Barbara, California, USA, April 2006.
- "Challenges and Opportunities in an Internet with Pervasive Wireless Access," University of Texas--Dallas Computer Science Colloquium, Dallas, Texas, USA, March 2006.
- "Challenges and Opportunities with Pervasive Wireless in the Internet," Duke University Computer Science Colloquium, Durham, North Carolina, USA, February 2006.
- "The Span From Wireless Protocols to Social Applications," Intel Research Labs, Cambridge, United Kingdom, December 2005.
- "The Internet Dot.Com Bomb and Beyond the Dot.Com Calm," CSE IGERT and Cal Poly Lecture Series, San Luis Obispo, California, USA, October 2005.
- "Panel: Directions in Networking Research," IEEE Computer Communications Workshop (CCW), Irvine, California, USA, October 2005.
- "Economic Incentives for Ad Hoc Networks," KAIST New Applications Seminar, Seoul, South Korea, March 2004.
- "New Applications for the Next Generation Internet," Citrix Systems, Santa Barbara, California, USA, March 2004.
- "PI: The Imperfect Pursuit of Pure Pattern," CITS Visions in Technology Series, Santa Barbara, California, USA, January 2004.
- "Panel: Core Networking Issues and Protocols for the Internet," National Science Foundation (NSF) Division of Advanced Networking Infrastructure and Research (ANIR) Principal Investigators Workshop, Washington DC, USA, March 2003.

- "Panel: Pricing for Content in the Internet," SPIE ITCom Internet Performance and Control of Network Systems, Boston, Massachusetts, USA, July 2002.
- "The Technology Behind Wireless LANs," Central Coast MIT Enterprise Forum, Santa Barbara, California, USA, March 2002.
- "Lessons Learned in the Digital Classroom," Center for Information and Technology Brown Bag Symposium, Santa Barbara, California, USA, March 2002.
- "The Evolution of Advanced Networking Services: From the ARPAnet to Internet2," California State University--San Luis Obispo CS Centennial Colloquium Series, San Luis Obispo, California, USA, February 2002.
- "Deployment of IP Multicast in Campus Infrastructures," Internet2 Campus Deployment Workshop, Atlanta, Georgia, USA, May 2001.
- "Multicast: Is There Anything Else to Do?," Sprint Research Retreat, Miami, Florida, USA, May 2001.
- "The Evolution of Next-Generation Internet Services and Applications," Government Technology Conference 2001 (GTC) for the Western Region, Sacramento, California, USA, May 2001.
- "I2 Multicast: Not WIDE-scale Deployment, FULL-scale Deployment," Closing Plenary, Internet2 Member Meetings, Washington, D.C., USA, March 2001.
- "Panel: Beyond IP Multicast," Content Delivery Networks (CDN), New York, New York, USA, February 2001.
- "Viable Multicast Pricing & Business Models for Wider-Scale Deployment," Content Delivery Networks (CDN), New York, New York, USA, February 2001.
- "IP Multicast: Modern Protocols, Deployment, and Management," Content Delivery Networks (CDN), New York, New York, USA, February 2001 & San Jose, California, USA, December 2001.
- "Under the Hood of the Internet," Technology 101: Technology for Investors, Center for Entrepreneurship & Engineering Management, November 2000.
- "Understanding Multicast Traffic in the Internet," (1) University of Virginia, (2) University of Maryland, and (3) Columbia University, September 2000.
- "The Bad, The Ugly, and The Good: The Past, Present, and Future of Multicast," Digital Fountain, San Francisco, California, USA, August 2000.
- "Implications of Source-Specific Multicast (SSM) on the Future of Internet Content Delivery," Occam Networks, Santa Barbara, California, USA, August 2000.
- "Introduction to Multicast Routing Protocols," UC-Berkeley Open Mesh Multicast Workshop, Berkeley, California, USA, July 2000.
- "Efforts to Understand Traffic and Tree Characteristics," University of Massachusetts--Amherst Colloquia, Amherst, Massachusetts, USA, May 2000.
- "Monitoring Multicast Traffic," Sprint Research Retreat, Half Moon Bay, California, USA, April 2000.
- "What is the Next Generation of Multicast in the Internet?," HRL Laboratories, Malibu, California,

USA, January 2000.

- "Mission and Status of the Center for Information Technology and Society (CITS)," Intel Research Council, Portland, Oregon, USA, September 1999.
- "Multicast at a Crossroads," IP Multicast Initiative Summits and Bandwidth Management Workshops, San Francisco, CA, USA, (1) October 1999; (2) February 2000; and (3) June 2000.
- "IP Multicast: Modern Protocols, Deployment, and Management," Networld+Interop: (1) Las Vegas, Nevada, USA--May 2000; (2) Tokyo, JAPAN--June 2000; (3) Atlanta, Georgia, USA--September 2000; (4) Las Vegas, Nevada, USA--May 2001; (5) Las Vegas, Nevada, USA--May 2002.
- "IP Multicast: Practice and Theory" (w/ Steve Deering), Networld+Interop: (1) Las Vegas, Nevada, USA--May 1999; (2) Tokyo, JAPAN--June 1999; and (3) Atlanta, Georgia, USA--September 1999.
- "Internet2 Multicast Testbeds and Applications," Workshop on Protocols for High Speed Networks (PfHSN), Salem, Massachusetts, USA, August 1999.
- "IP Multicast: Protocols for the Intra- and Inter-Domain," Lucent Technologies, Westford, Massachusetts, USA, August 1999.
- "Internet2 Multicast Testbeds and Applications," NASA Workshop: Bridging the Gap, Moffett Field, California, USA, August 1999.
- "The Evolution of Next-Generation Services and Applications in the Internet," Tektronix Distinguished Lecture Series, Portland, Oregon, USA, May 1999.
- "Multicast Applications and Infrastructure in the Next Generation Internet," CENIC 99 Workshop on Achieving Critical Mass for Advanced Applications, Monterey, California, USA, May 1999.
- "Multicast Traffic Monitoring and Analysis Work at UCSB" (w/ P. Rajvaidya), Workshop on Internet Statistics and Metrics Analysis (ISMA), San Diego, California, USA, April 1999.
- "How the Internet Works: Following Bits Around the World," Science Lite, Santa Barbara General Affiliates and Office of Community Relations, California, USA, February 1999.
- "Managing Multicast: Challenges, Tools, and the Future," IP Multicast Initiative Summit, San Jose, California, USA, February 1999.
- "The Future of Multicast Communication and Protocols," Internet Bandwidth Management Summit (iBAND), San Jose, California, USA, November 1998.
- "An Overview of IP Multicast: Applications and Deployment," (1) Workshop on Evaluating IP Multicast as the Solution for Webcasting Real-Time Networked Multimedia Applications, New York, New York, USA, July 1998; and (2) Satellites and the Internet Conference, Washington, D.C., USA, July 1998.
- "IETF Developments in IP Multicast," IP Multicast Initiative Summit, San Jose, California, USA, February 1998.
- "An Introduction to IP Multicast and the Multicast Backbone (MBone)" vBNS Technical Meeting sponsored by the National Center for Network Engineering (NLANR), San Diego, California, USA, February 1998.

- "Using Multicast Communication to Deliver WWW Pages" Computer Communications Workshop (CCW '97), Phoenix, Arizona, USA, September 1997.

G. Research Funding

- K. Almeroth, "Packet Scheduling Using IP Embedded Transport Instrumentation," Cisco Systems Inc., \$100,000, 3/1/13-8/31/14.
- K. Almeroth, E. Belding and S.J. Lee, "GOALI: Maximizing Available Bandwidth in Next Generation WLANs", National Science Foundation (NSF), \$101,088, 10/1/13-9/30/14.
- K. Almeroth and E. Belding, "GOALI: Intelligent Channel Management in 802.11n Networks," National Science Foundation (NSF), \$51,000, 10/1/10-9/30/11.
- B. Zhao, K. Almeroth, H. Zheng, and E. Belding, "NeTS: Medium: Airlab: Distributed Infrastructure for Wireless Measurements," National Science Foundation (NSF), \$700,000, 9/1/09-8/13/13.
- K. Almeroth, E. Belding and T. Hollerer, "NeTS-WN: Wireless Network Health: Real-Time Diagnosis, Adaptation, and Management," National Science Foundation (NSF), \$600,000, 10/1/07-9/30/10.
- K. Almeroth, "Next-Generation Service Engineering in Internet2," University Consortium for Advanced Internet Development (UCAID), \$1,254,000, 7/1/04-6/30/09 (reviewed and renewed each year).
- B. Manjunath, K. Almeroth, F. Bullo, J. Hespanha, T. Hollerer, C. Krintz, U. Madhow, K. Rose, A. Singh, and M. Turk, "Large-Scale Multimodal Wireless Sensor Network," Office of Naval Research Defense University Research Instrumentation Program (DURIP), \$655,174, 4/14/08-4/14/09.
- K. Almeroth and E. Belding, "Improving Robustness in Evolving Wireless Infrastructures," Intel Corporation, \$135,000, 7/1/06-6/30/09 (reviewed and renewed for second and third year).
- K. Almeroth and K. Sarac, "Bridging Support in Mixed Deployment Multicast Environments," Cisco Systems Inc., \$100,000, 9/1/07-8/31/08.
- K. Sarac and K. Almeroth, "Building the Final Piece in One-to-Many Content Distribution," Cisco Systems Inc., \$95,000, 9/1/06-8/31/07.
- E. Belding, K. Almeroth and J. Gibson, "Real-Time Communication Support in a Ubiquitous Next-Generation Internet," National Science Foundation (NSF), \$900,000, 10/1/04-9/30/07.
- K. Almeroth and K. Sarac, "Improving the Robustness of Multicast in the Internet," Cisco Systems Inc., \$80,000, 9/1/04-8/31/05.
- R. Mayer, B. Bimber, K. Almeroth and D. Chun, "Assessing the Pedagogical Implications of Technology in College Courses," Mellon Foundation, \$350,000, 7/1/04-6/30/07.
- B. Bimber, A. Flanagin and C. Stol, "Technological Change and Collective Association: Changing Relationships Among Technology, Organizations, Society and the Citizenry," National Science Foundation (NSF), \$329,175, 7/1/04-6/30/07.
- K. Almeroth and B. Bimber, "Plagiarism Detection Techniques and Software," UCSB Instructional

Development, \$22,000, 7/1/04-6/30/05.

- K. Almeroth, "Student Travel Support for the 14th International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV)," National Science Foundation (NSF), \$10,000, 5/1/04-8/31/04.
- K. Almeroth, "An Automated Indexing System for Remote, Archived Presentations," QAD Inc., \$25,000, 5/1/04-6/30/05.
- K. Almeroth and M. Turk, "A Remote Teaching Assistant Support System," Microsoft, \$40,000, 1/1/04-6/30/05.
- K. Almeroth, "Supporting Multicast Service Functionality in Helix," Real Networks, \$30,000, 9/1/03-6/30/04.
- K. Almeroth and E. Belding, "Service Discovery in Mobile Networks," Nokia Summer Research Grant (U. Mohan), \$10,240, 7/1/03-9/30/03.
- K. Almeroth, D. Zappala, "Building a Global Multicast Service," Cisco Systems Inc., \$100,000, 1/1/03-6/30/04.
- K. Almeroth, "Developing A Dynamic Protocol for Candidate Access Router Discovery," Nokia Graduate Student Fellowship (R. Chalmers), \$26,110, 9/01/02-6/30/03.
- B. Bimber and K. Almeroth, "The Role of Collaborative Groupware in Organizations," Toole Family Foundation, \$182,500 (\$20,000 cash plus \$162,500 in software), 9/1/02-8/30/07.
- B. Manjunath, et al., "Digital Multimedia: Graduate Training Program in Interactive Digital Multimedia," National Science Foundation (NSF), \$2,629,373, 4/1/02-3/31/07.
- J. Green, K. Almeroth, et al., "Inquiry in the Online Context: Learning from the Past, Informing the Future," UCSB Research Across Disciplines, \$10,000, 9/1/01-8/31/02.
- K. Almeroth, "Monitoring and Maintaining the Global Multicast Infrastructure," Cisco Systems Inc., \$54,600, 7/1/01-6/30/02.
- R. Kemmerer, K. Almeroth, et al., "Hi-DRA High-speed, Wide-area Network Detection, Response, and Analysis," Department of Defense (DoD), \$4,283,500, 5/1/01-4/30/06.
- A. Singh, K. Almeroth, et al., "Digital Campus: Scalable Information Services on a Campus-wide Wireless Network," National Science Foundation (NSF), 1,450,000, 9/15/00-12/31/04.
- K. Almeroth, "Visualizing the Global Multicast Infrastructure," UC MICRO w/ Cisco Systems Inc., \$85,438, 7/1/00-6/30/02.
- H. Lee, K. Almeroth, et al., "Dynamic Sensing Systems," International Telemetering Foundation, \$260,000, 07/01/00-06/30/04.
- B. Bimber and K. Almeroth, "Funding for the Center on Information Technology and Society," \$250,000 from Dialogic (an Intel Company) and \$250,000 from Canadian Pacific.
- K. Almeroth, "CAREER: From Protocol Support to Applications: Elevating Multicast to a Ubiquitous Network Service," National Science Foundation (NSF), \$200,000, 9/1/00-8/31/04.

- K. Almeroth, "Characterizing Multicast Use and Efficiency in the Inter-Domain," Sprint Advanced Technology Laboratories, \$62,500, 3/1/00-6/30/01.
- K. Almeroth, "Producing the Next Generation of Multicast Monitoring and Management Protocols and Tools," UC MICRO w/ Cisco Systems Inc., \$124,500, 7/1/99 - 6/30/01.
- K. Almeroth, "Utilizing Satellite Links in the Provision of an Inter-Wide Multicast Service," HRL Laboratories, \$20,000, 7/1/99 - 6/30/00.
- T. Smith, K. Almeroth, et al., "Alexandria Digital Earth Prototype," National Science Foundation, \$5,400,000, 4/1/99-3/31/04.
- V. Vesna, K. Almeroth, et al., "Online Public Spaces: Multidisciplinary Explorations in Multi-User Environments (OPS:MEME), Phase II," UCSB Research Across Disciplines, \$50,000, 9/1/98-8/31/99.
- K. Almeroth, "Techniques and Analysis for the Provision of Multicast Route Management," UC MICRO w/ Cisco Systems Inc., \$97,610, 7/1/98 - 6/30/00.
- K. Almeroth, "Capturing and Modeling Multicast Group Membership in the Multicast Backbone (MBone)," UC MICRO w/ Hughes Research Labs, \$19,146, 7/1/98 - 12/31/99.
- K. Almeroth, "Building a Content Server for the Next Generation Digital Classroom," UCSB Faculty Research Grant, \$5,000, 7/1/98-6/31/99.

H. Research Honors and Awards

- IEEE Fellow Status, 2013
- Finalist for Best Paper Award, IEEE Conference on Sensor and Ad Hoc Communications and Networks (SECON), June 2008
- Best Paper Award, Passive and Active Measurement (PAM) Conference, April 2007
- Outstanding Paper Award, World Conference on Educational Multimedia, Hypermedia & Telecommunications (ED MEDIA), June 2006
- IEEE Senior Member Status, 2003
- Finalist for Best Student Paper Award, ACM Multimedia, December 2002
- Outstanding Paper Award, World Conference on Educational Multimedia, Hypermedia & Telecommunications (ED MEDIA), June 2002
- Computing Research Association (CRA) Digital Government Fellowship, 2001
- National Science Foundation CAREER Award, 2000
- Best Paper Award, 7th International World Wide Web Conference, April 1998

III. Service

A. Professional Activities

1. Society Memberships

Member, Association for Computing Machinery (ACM): 1993-present
 Member, ACM Special Interest Group on Communications (SIGComm): 1993-present
 Fellow, Institute of Electrical and Electronics Engineers (IEEE): 1993-present
 Member, IEEE Communications Society (IEEE ComSoc): 1993-present
 Member, American Society for Engineering Education (ASEE): 2003-2006

2. Review Work for Technical Journals and Publishers

NSF CISE research proposals, IEEE/ACM Transactions on Networking, IEEE/ACM Transactions on Computers, IEEE/ACM Transactions on Communications, IEEE Transactions on Circuits and Systems for Video Technology, IEEE Transactions on Parallel and Distributed Systems, IEEE Transactions on Multimedia, IEEE Communications, IEEE Communications Letters, IEEE Network, IEEE Internet Computing, IEEE Multimedia, IEEE Aerospace & Electronics Systems Magazine, ACM Transactions on Internet Technology, ACM Transactions on Multimedia Computing, Communications and Applications, ACM Computing Surveys, ACM Computer Communications Review, ACM Computeres in Entertainment, ACM/Springer Multimedia Systems Journal, AACE Journal of Interactive Learning (JILR), International Journal of Computer Mathematics, Journal of Communications and Networks, Journal of Parallel and Distributed Computing, Journal of Network and Systems Management, Journal of High Speed Networking, Journal of Communications and Networks, Journal on Selected Areas in Communications, Journal of Wireless Personal Communications, Personal Mobile Communications, Annals of Telecommunications, International Journal of Wireless and Mobile Computing, Pervasive and Mobile Computing (PMC), Wireless Networks Journal, Computer Networks Journal, Cluster Computing, Computer Communications, Mobile Computing and Communications Review, Performance Evaluation, Software--Practice & Experience, Information Processing Letters, ACM Sigcomm, ACM Multimedia, ACM Network and System Support for Digital Audio and Video Workshop (NOSSDAV), ACM Sigcomm Workshop on the Economics of Peer-to-Peer Systems (P2PEcon), ACM Sigcomm Workshop on Challenged Networks (CHANTS), IEEE Infocom, IEEE Globecom, IEEE Global Internet (GI) Symposium, IEEE Globecom Automatic Internet Symposium, IEEE Globecom Internet Services and Enabling Technologies (IS&ET) Symposium, IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM), IEEE International Conference on Network Protocols (ICNP), IEEE Conference on Sensor and Ad Hoc Communications and Networks (SECON), IEEE International Conference on Multimedia and Exposition (ICME), IEEE International Conference on Communications (ICC), IEEE International Conference on Parallel and Distributed Systems (ICPADS) IEEE International Symposium on High-Performance Distributed Computing (HPDC), IEEE International Conference on Distributed Computing Systems (ICDCS), IEEE International Workshop on Quality of Service (IWQoS), IEEE/IFIP Network Operations and Management Symposium (NOMS), IFIP/IEEE International Symposium on Integrated Network Management (IM), IFIP/IEEE International Conference on Management of Multimedia Networks and Services (MMNS), IEEE Aerospace & Electronics Systems Magazine, SPIE Conference on Multimedia Computing and Networking (MMCN), IFIP Networking, IASTED International Conference on Information Systems and Databases (ISD), IASTED International Conference on Communications, Internet, and Information Technology, IASTED International Conference on Internet and Multimedia Systems and Applications (IMSA), IASTED International Conference on European Internet and Multimedia Systems and Applications (EuroIMSA), IASTED International Conference on Communications and Computer Networks (CCN), IASTED International Conference on Software Engineering and Applications (SEA), International Conference on Computer and Information Science (ICIS),

International Association for Development of the Information Society (IADIS) International Conference on the WWW/Internet, Workshop on Network Group Communication (NGC), International Conference on Next Generation Communication (CoNEXT), International Conference on Parallel Processing (ICPP), International Conference on Computer Communications and Networks (IC3N), International Workshop on Hot Topics in Peer-to-Peer Systems (Hot-P2P), International Workshop on Wireless Network Measurements (WiNMee), International Workshop on Incentive-Based Computing (IBC), International Workshop on Multi-hop Ad Hoc Networks (REALMAN), International Workshop on Broadband Wireless Multimedia: Algorithms, Architectures and Applications (BroadWIM), International Packet Video (PV) Workshop, High Performance Networking Conference (HPN), International Parallel Processing Symposium (IPPS), International Symposium on Innovation in Information & Communication Technology (ISIICT), Workshop on Coordinated Quality of Service in Distributed Systems (COQODS), Pearson Education (Cisco Press) Publishers, Macmillan Technical Publishing, and Prentice Hall Publishers.

3. Conference Committee Activities

Journal/Magazine Editorial Board

IEEE Transactions on Mobile Computing (TMC): 2006-2011, 2017-2020 (Associate Editor-in-Chief)
 IEEE Networking Letters: 2018-present
 IEEE Transactions on Network and Service Management (TNSM): 2015-present
 Journal of Network and Systems Management (JNSM): 2011-present
 IEEE/ACM Transactions on Networking (ToN): 2003-2009, 2013-2017
 ACM Computers in Entertainment: 2002-2015
 IEEE Network: 1999-2012
 AACE Journal of Interactive Learning Research (JILR): 2003-2012
 IEEE Transactions on Mobile Computing (TMC): 2006-2011
 ACM Computer Communications Review (CCR): 2006-2010

Journal/Magazine Guest Editorship

IEEE Journal on Selected Areas in Communications (JSAC) Special Issue on "Delay and Disruption Tolerant Wireless Communication", June 2008
 Computer Communications Special Issue on "Monitoring and Measuring IP Networks", Summer 2005
 Computer Communications Special Issue on "Integrating Multicast into the Internet", March 2001

Conference/Workshop Steering Committee

IEEE International Conference on Network Protocols (ICNP): 2007-present
 ACM Sigcomm Workshop on Challenged Networks (CHANTS): 2006-present
 International Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV): 2001-present, 2005-2011 (chair), 2012-2020 (co-chair)
 IEEE Global Internet (GI) Symposium: 2005-2013, 2018-present
 IFIP/IEEE International Conference on Management of Multimedia Networks and Services (MMNS): 2005-2009

Conference/Workshop Chair

International Conference on Communication Systems and Networks (COMSNETS): 2014 (co-chair)
ACM International Conference on Next Generation Communication (CoNext): 2013 (co-chair)
ACM RecSys News Recommender Systems (NRS) Workshop and Challange: 2013 (co-chair)
ACM Sigcomm Workshop on Challenged Networks (CHANTS): 2006 (co-chair)
IEEE International Conference on Network Protocols (ICNP): 2003 (co-chair), 2006
International Workshop on Wireless Network Measurements (WiNMee): 2006 (co-chair)
IFIP/IEEE International Conference on Management of Multimedia Networks and Services (MMNS): 2002 (co-chair)
International Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV): 2002 (co-chair), 2003 (co-chair)
IEEE Global Internet (GI) Symposium: 2001 (co-chair), 2018 (co-chair)
International Workshop on Networked Group Communication (NGC): 2000 (co-chair)

Program Chair

International Conference on Computer Communication and Networks (ICCCN): 2015
(Track co-chair) International Conference on Communication Systems and Networks (COMSNETS): 2010
IEEE International Conference on Network Protocols (ICNP): 2008 (co-chair)
IEEE Conference on Sensor and Ad Hoc Communications and Networks (SECON): 2007 (co-chair)
IFIP Networking: 2005 (co-chair)

Posters/Demonstrations Chair

ACM Sigcomm: 2012 (co-chair)

Student Travel Grants Chair

ACM Sigcomm: 2010 (co-chair)

Publicity Chair

IFIP/IEEE International Conference on Management of Multimedia Networks and Services (MMNS): 2004 (co-chair)

Keynote Chair

IEEE Infocom: 2005 (co-chair)

Local Arrangements Chair

Internet2 "Field of Dreams" Workshop: 2000

Tutorial Chair

ACM Multimedia: 2000

IEEE International Conference on Network Protocols (ICNP): 1999

Panel/Session Organizer

NSF ANIR PI 2003 Panel on "Core Networking Issues and Protocols for the Internet"
 CCW 2001 Session on "Multicast/Peer-to-Peer Networking"
 NOSSDAV 2001 Panel on "Multimedia After a Decade of Research"
 NGC 2000 Panel on "Multicast Pricing"

Technical Program Committee

IEEE International Conference on Network Protocols (ICNP): 1999, 2000, 2001, 2003, 2004, 2005, 2006, 2007, 2008, 2009 (Area Chair), 2010 (Area Chair), 2011 (Area Chair), 2012 (Area Chair), 2013, 2014 (Area Chair), 2015 (Area Chair), 2016 (Area Chair), 2017 (Area Chair), 2018 (Area Chair), 2019 (Area Chair)
 International Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV): 2000, 2001, 2002, 2003, 2004, 2005, 2006, 2007, 2008, 2009, 2010, 2011, 2012, 2013, 2014, 2015, 2016, 2017, 2018, 2019
 ACM Multimedia (MM): 2001, 2003, 2004, 2005 (short paper), 2006, 2007, 2008, 2008 (short paper), 2010, 2011, 2012, 2013, 2015, 2016, 2017, 2018, 2019
 IEEE Conference on Sensor and Ad Hoc Communications and Networks (SECON): 2004, 2005, 2006, 2007, 2008, 2009, 2010, 2011 (Area Chair), 2012 (Area Chair), 2013, 2014 (Area Chair), 2015, 2016 (Area Chair), 2017, 2018, 2019
 IEEE/IFIP Network Operations and Management Symposium (NOMS): 2004, 2006, 2010
 IEEE Infocom: 2004, 2005, 2006, 2008, 2009, 2010 (Area Chair), 2011 (Area Chair), 2012 (Area Chair)
 IFIP Networking: 2004, 2005, 2006, 2007, 2010, 2011, 2012, 2013, 2014, 2015, 2016, 2017, 2018, 2019
 IEEE International Conference on Communications (ICC) Next Generation Networking and Internet Symposium (NGNI): 2018, 2019
 ACM Workshop on Mobile Video (MoVid): 2014, 2015, 2016, 2017
 ACM Student Research Competition (SRC) Grand Finals: 2014
 Mobile and Social Computing for Collaborative Interactions (MSC): 2014
 IEEE Conference on Communications and Network Security (CNS): 2013
 IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM): 2005, 2006, 2007, 2008, 2009, 2010
 ACM Sigcomm Workshop on Challenged Networks (CHANTS): 2006, 2008, 2009, 2010, 2011, 2012, 2016, 2017, 2018, 2019
 IEEE International Conference on Distributed Computing Systems (ICDCS): 2006, 2010, 2011, 2012, 2013
 International Workshop on Wireless Network Measurements (WiNMee): 2006, 2008, 2010
 ACM Sigcomm: 2008 (poster), 2010
 IEEE International Conference on Computer Communication and Networks (IC3N): 2008, 2009, 2010, 2011, 2012
 International Conference on Communication Systems and Networks (COMSNETS): 2009, 2010, 2011, 2012, 2013
 International Conference on Sensor Networks (SENSORNETS): 2012
 International Workshop on Social and Mobile Computing for Collaborative Environments (SOMOCO): 2012
 Workshop on Scenarios for Network Evaluation Studies (SCENES): 2009, 2010, 2011
 ACM Multimedia Systems (MMSys): 2010, 2011, 2012, 2015, 2016, 2017, 2018, 2019
 IEEE International Symposium on Multimedia (ISM): 2016
 IEEE International Conference on Pervasive Computing and Communications (PerCom): 2010
 IEEE Wireless Communications and Networking Conference (WCNC): 2010, 2011

ACM International Symposium on Mobility Management and Wireless Access (MobiWac): 2010, 2011

International Conference on Computing, Networking and Communications, Internet Services and Applications Symposium (ICNC-ISA): 2012, 2013

IEEE WoWMoM Workshop on Hot Topics in Mesh Networking (HotMesh): 2010, 2011, 2012, 2013

IEEE Workshop on Pervasive Group Communication (PerGroup): 2010

ACM International Conference on Next Generation Communication (CoNEXT): 2005, 2006, 2007, 2009, 2012

IEEE International Conference on Broadband Communications, Networks, and Systems (BroadNets) Wireless Communications, Networks and Systems Symposium: 2007, 2008, 2009

IEEE International Conference on Broadband Communications, Networks, and Systems (BroadNets) Internet Technologies Symposium: 2007, 2008, 2009

International Workshop on Mobile and Networking Technologies for Social Applications (MONET): 2008, 2009

Extreme Workshop on Communication-The Midnight Sun Expedition (ExtremeCom): 2009

IEEE International Workshop on Cooperation in Pervasive Environments (CoPE): 2009

International Workshop on the Network of the Future (FutureNet): 2009, 2010, 2011, 2012

IEEE International Conference on Multimedia and Exposition (ICME): 2010

SPIE Conference on Multimedia Computing and Networking (MMCN): 2004, 2008

ACM Sigcomm Workshop on the Economics of Networks, Systems, and Computation (NetEcon): 2008

IEEE International Conference on Communications (ICC): 2008

IEEE International Conference on Mobile Ad-hoc and Sensor Systems (MASS): 2008

IFIP/IEEE International Symposium on Integrated Network Management (IM): 2005, 2007

Global Internet (GI) Symposium: 2001, 2002, 2004, 2006, 2007

IEEE/ACM International Conference on High Performance Computing (HiPC): 2007

ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc): 2007

IEEE Workshop on Embedded Systems for Real-Time Multimedia (ESTIMedia): 2007

IEEE/IFIP Wireless On Demand Network Systems and Services (WONS): 2007

IFIP/IEEE International Conference on Management of Multimedia Networks and Services (MMNS): 2001, 2002, 2003, 2004, 2005, 2006

IASTED International Conference on European Internet and Multimedia Systems and Applications (EuroIMSA): 2004, 2006

IEEE International Conference on Parallel and Distributed Systems (ICPADS): 2005, 2006

IEEE Globecom Internet Services and Enabling Technologies (IS&ET) Symposium: 2006

International Workshop on Incentive-Based Computing (IBC): 2006

IEEE International Workshop on Quality of Service (IWQoS): 2006, 2014, 2015

International Workshop on Multi-hop Ad Hoc Networks (REALMAN): 2006

IEEE Globecom Automatic Internet Symposium: 2005

ACM Sigcomm Workshop on the Economics of Peer-to-Peer Systems (P2PEcon): 2005

International Conference on Parallel Processing (ICPP): 2001, 2003, 2004

International Packet Video (PV) Workshop: 2002, 2003, 2004

IEEE International Symposium on High-Performance Distributed Computing (HPDC): 2004

ACM Sigcomm: 2004 (poster)

International Workshop on Broadband Wireless Multimedia: Algorithms, Architectures

and Applications (BroadWIM): 2004
International Symposium on Innovation in Information & Communication Technology (ISIICT): 2004
Workshop on Coordinated Quality of Service in Distributed Systems (COQODS): 2004
IASTED International Conference on Networks and Communication Systems (NCS): 2004
IASTED International Conference on Communications, Internet, and Information Technology (CIIT): 2004
IASTED International Conference on Internet and Multimedia Systems and Applications (IMSA): 2003, 2004
International Workshop on Networked Group Communication (NGC): 1999, 2000, 2001, 2002, 2003
International Association for Development of the Information Society (IADIS)
International Conference on the WWW/Internet: 2003
International Conference on Computer and Information Science (ICIS): 2003
Human.Society@Internet: 2003
IASTED International Conference on Communications and Computer Networks (CCN): 2002
The Content Delivery Networks (CDN) Event: 2001
IP Multicast Initiative Summit: 1998, 1999, 2000
Corporation for Education Network Initiatives in California (CENIC): 1999
Internet Bandwidth Management Summit (iBAND): 1998, 1999

B. Technical Activities

1. Working Groups

Internet2 Working Group on Multicast, Chair: 1998-2005
IEEE Communications Society Internet Technical Committee (ITC), Conference Coordinator: 2000-2004
IETF Multicast Directorate (MADDOGS), Member: 1999-2001
IASTED Technical Committee on the Web, Internet and Multimedia, Member: 2002-2005
Internet Engineering Task Force (IETF), various working groups: 1995-present

2. Meeting Support Work

Internet Engineering Task Force MBone broadcasts: 1995-2005
Conference MBone broadcasts: Sigcomm '99, and '00
Interop+Networld Network Operations Center (NOC) Team Member: 1995-1997
ACM Multimedia technical staff: 1994

C. University of California Committees

1. Department of Computer Science Committees

Public Relations: 2005-2006 (chair 2005-2006), 2009-2011 (chair 2009-2011)
Strategic Planning: 2000-2002, 2003-2006, 2009-2011
Undergraduate Advising and Affairs: 2006-2007, 2014-2015
Vice Chair: 2000-2005
Graduate Admissions: 2000-2005 (chair 2000-2005), 2011-2012
Graduate Affairs: 2000-2005 (co-chair 2000-2005)
Teaching Administration: 2000-2005
Facilities: 1997-2001 (chair 1999-2000), 2006-2007
External Relations: 1999-2002
Computer Engineering Public Relations: 2011-2012
Computer Engineering Awards: 2011-2012
Computer Engineering Administration/Recruiting: 1998-2001
Computer Engineering Lab and Computer Support: 1998-2001
Faculty Recruiting: 1999-2002
Graduate Advising: 1998-1999, 2000-2005

2. University Committees

Member, Campus Budget and Planning: 2013-2015
Faculty, Cognitive Science Program: 2006-present
Faculty, Technology Management Program (TMP): 2003-2014
Faculty, Media Arts and Technology (MAT) Program: 1998-2014
Faculty, Computer Engineering Degree Program: 1998-present
Steering Committee, Center for Information Technology and Society (CITS): 2012-present
Associate Director, Center for Information Technology and Society (CITS): 1999-2012
Member, Campus Committee on Committees: 2010-2013
Member, Campus Income and Recharge Committee: 2010-2013
Member, College of Engineering Executive Committee: 2010-2012 (chair 2011-2012), 2014-2015 (chair 2014-2015)
Member, Distinguished Teaching Award Committee: 2009, 2010, 2011
Member, Campus Classroom Design and Renovation Committee: 2003-2010
Member, ISBER Advisory Committee: 2008-2011
Member, Fulbright Campus Review Committee: 2007
Member, Faculty Outreach Grant Program Review Committee: 2007
Executive Vice Chancellor's Information Technology Fee Committee: 2005-2006
Council on Research and Instructional Resources: 2003-2006
Executive Vice Chancellor's Working Group on Graduate Diversity: 2004-2005
Member, Engineering Pavilion Planning Committee: 2003-2005
Information Technology Board: 2001-2004
Executive Committee, Center for Entrepreneurship & Engineering Management (CEEM): 2001-2004

3. System Wide Committees

UCSB Representative to the Committee on Information Technology and Telecommunications Policy (ITTP): 2003-2005
UCSB Representative to the Executive Committee, Digital Media Innovation (DiMI): 1998-2003

D. Georgia Tech Committees and Service (while a graduate student)

Graduate Student Body President: 1994-1995

Georgia Tech Executive Board: 1994-1995

Georgia Tech Alumni Association Executive Committee: 1994-1995

Dean of Students National Search Committee: 1995

Institute Strategic Planning Committee: 1994-1996

Cases in last 4 years I have been deposed or testified:

- A deposition in Inter Partes Review of U.S. Patent No. 7,516,177 (IPR2016-01434) [Oracle America, Inc. and Oracle Corporation and HCC Insurance Holdings, Inc. v. Intellectual Ventures II]. Finished 07/17.
- A deposition in Comcast Cable Communications, LLC v. OpenTV, Inc. and Nagravision SA (3:16-CV-6180-WHA, N.D. Cal.). Finished 07/17.
- Depositions in Thomas C. Sisoian v. International Business Machines Corporation (A-14-CA-565-SS, W.D. Tex.). Finished 07/17.
- A deposition in Inter Partes Review of U.S. Patent Nos. 6,895,449 and 6,470,399 (IPR2017-00713 and IPR2017-00714, respectively) [ZTE (USA) Inc. v. Papst Licensing GmbH & Co. KG.]. Finished 08/17.
- Depositions and trial testimony in Packet Intelligence, LLC v. NetScout Systems, Inc. and Sandvine Corporation (2:16-CV-00230, 2:16-CV-00147, E.D. Tex.). Finished 11/17.
- A deposition in in Certain Two-Way Radio Equipment and Systems, Related Software and Components Thereof (US ITC Inv. Nos. 337-TA-1053) [Motorola v. Hytera]. Finished 11/17.
- A deposition in Sound View Innovations, LLC v. Facebook, Inc. (16-116-RGA, D. Del). Finished 11/17.
- A deposition in Sprint Communications Company LP v. Cox Communications, Inc. (12-487-SLR, D. Del.). Finished 12/17.
- Depositions in Alacritech, Inc. v. Centurylink Communications LLC; Winstron Corporation, Dell, Inc. (2:16-cv-693-RWS, 2:16-cv-692-RWS, 2:16-cv-695-RWS, E.D. Tex.). Finished 12/17.
- A deposition and trial testimony in Sonos, Inc. v. D&M Holding, et al. (14-1330-RGA, D. Del.). Finished 12/17.
- Depositions in Inter Partes Review of U.S. Patent Nos. 7,899,492; 8,050,711; 9,355,611; 8,712,471; 9,286,853; 8,903,451; 8,948,814 and 9,118,794 (IPR2017-00870 through IPR2017-00879) [HTC America, Inc. v. Virginia Innovation Sciences, Inc.]. Finished 01/18.
- Depositions and trial testimony in Certain Network Devices, Related Software and Components Thereof (II) (US ITC Inv. Nos. 337-TA-945 and 337-TA-945M) [Cisco Systems, Inc. v. Arista Networks, Inc.]. Finished 01/18.
- Depositions in Inter Partes Review of U.S. Patent Nos. 8,000,314 (IPR2015-01901); 8,013,732 (IPR2015-01973); 8,754,780 (IPR2016-00984); 8,908,842 (CBM2016-00095); 6,249,516 (IPR2016-01602); 7,697,492 (IPR2016-01895); 7,468,661 (IPR2017-00001); 8,233,471 (IPR2017-00007 and IPR2017-00008); 8,625,496 (IPR2017-00213); 8,013,732 (IPR2017-00216); 6,437,692 (IPR2017-00359); and 8,000,314 (IPR2017-00252) [Emerson Electric Co. v. (S)IPCO, LLC]. Finished 01/18.
- A deposition in D&M Holding, Inc., et al. v. Sonos, Inc. v. (16-141-RGA, D. Del.). Finished 03/18.

- A deposition in Limelight Networks, Inc. v. XO Communication, LLC, Akamai Technologies, Inc., and MIT (3:15cv720-JAG, E.D. Va.). Finished 03/18.
- A deposition in Twilio, Inc. v. Telesign Corporation (5:16-CV-06925-LHK-SVK, N.D. Cal.). Finished 04/18.
- Depositions in Arista Networks, Inc. v. Cisco Systems, Inc. (5:16-cv-00923-BLF, N.D. Cal.). Finished 08/18.
- Depositions in Inter Partes Review of U.S. Patent No. 9,083,997 (IPR2017-00829 and IPR2017-00830) [Twitter, Inc. v. YouToo Technologies, LLC]. Finished 09/18.
- A deposition in Richard A. Williamson, on behalf of and as a trustee for At Home Bondholders' Liquidating Trust v. Google, LLC (5:15-CV-00966-BLF, N.D. Cal.). Finished 10/18.
- A deposition in Implicit, LLC. v. Palo Alto Networks, Inc. (6:17-cv-182-JRG, E.D. Tex.). Finished 10/18.
- A deposition in Inter Partes Review of U.S. Patent Nos. 5,822,523 (IPR2018-00129 and IPR2018-00130) and 6,226,686 (IPR2018-00131 and IPR2018-00132) [Riot Games, Inc. v. Paltalk Holdings, Inc.]. Finished 11/18.
- Depositions in Inter Partes Review of U.S. Patent Nos. 8,116,284 (IPR2018-00128) and 6,591,111 (IPR2018-00176) [Hytera Communications Corp. LTD v. Motorola Solutions, Inc.]. Finished 01/19.
- Depositions in Inter Partes Review of U.S. Patent Nos. 7,237,036 (IPR2017-01391); 7,337,241 (IPR2017-01392); 9,055,104 (IPR2017-01393); 7,124,205 (IPR2017-01405 and IPR2018-00226); 7,673,072 (IPR2017-01406); 8,131,880 (IPR2017-1409 and IPR2017-1410); 8,805,948 (IPR2018-00234); and 7,945,699 (IPR2018-00401) [Intel Corp v. Alacritech, Inc.]. Finished 01/19.
- A claim construction tutorial in Ipsilium, LLC v. Cisco Systems, Inc. (4:17-cv-07179-HSG, N.D. Cal.). Finished 03/19.
- Depositions in Inter Partes Review of U.S. Patent Nos. 8,886,739 (IPR2018-00200); 9,306,885 (IPR2018-00312); 9,315,155 (IPR2018-00369); 9,306,886 (IPR2018-00397); 8,935,351 (IPR2018-00404); 9,338,111 (IPR2018-00408); 9,413,711 (IPR2018-00416 and IPR2018-00439); 9,313,157 (IPRs018-00455); and 9,313,156 (IPR2018-00458) [Snap, Inc. v. Vaporstream, Inc.]. Finished 04/19.
- Depositions in Rmail Limited, et al. v. Amazon.com, Inc.; Docusign, Inc.; Right Signature, LLC, et al.; and Adobe Systems Inc., et al. (2:10-cv-258-JRG, 2:11-cv-299-JRG, 2:11-cv-300-JRG, 2:11-cv-325-JRG, E.D. Tex.). Finished 06/19.
- Depositions and trial testimony in Certain Wireless Mesh Networking Products and Related Components Thereof (US ITC Inv. No. 337-TA-1131) [SIPCO, LLC v. Emerson Electric Co]. Finished 09/19.
- A deposition in Packet Intelligence v. Nokia of America Corporation (2:18-cv-382-JRG, E.D. Tex.). Finished 09/19.

- A deposition in Luminati Networks Ltd. v. UAB Tesonet and UAB Metacluster LT, (2:18-cv-00299-JRG, E.D. Tex.). Finished 11/19.
- A deposition in TQ Delta, LLC v. Zyxel Communications, Inc. (13-cv-2013-RGA, D. Del.). Finished 11/19.
- A deposition and trial testimony in Implicit, LLC v. NetScout Systems, Inc. and Sandvine Corporation (2:18-cv-0053-JRG and 2:18-cv-0054-JRG, E.D. Tex.). Finished 12/19.
- A deposition and trial testimony in Sony Music Entertainment, et al. v. Cox Communications, Inc. et al. (1:18-cv-00950-LO-JFA, E.D. Va.). Finished 12/19.
- Depositions in Inter Partes Review of U.S. Patent Nos. 9,369,741 (IPR2019-00231) and 9,621,956 (IPR2019-00281) [Comcast Cable Communications, LLC v. Rovi Guides Inc.]. Finished 01/20.
- A deposition and trial testimony in Certain Digital Video Receivers, Broadband Gateways, and Related Hardware and Software Components (US ITC Inv. No. 337-TA-1158) [Rovi Guides, Inc. v. Comcast Cable Communications, LLC]. Finished 01/20.
- A deposition in Innovation Sciences, LLC v. Resideo Technologies, Inc. (4:18-cv-00475-ALM, E.D. Tex.). Finished 02/20.
- Depositions and a claim construction tutorial in Blackberry Limited v. Facebook, Inc. and Snap, Inc. (2:18-cv-01844 GW(KSx) and 2:18-cv-02693-GW(KSx), C.D. Cal.). Finished 03/20.
- A deposition in Inter Partes Review of U.S. Patent Nos. 9,912,983 (IPR2019-01306) [Resideo Technologies, Inc. v. Innovation Sciences, LLC]. Finished 05/20.
- A deposition and trial testimony in Centripetal Networks, Inc. v. Cisco Systems, Inc. (2:18-CV-00094, E. D. Va.). Finished 05/20.
- Depositions in Inter Partes Review of U.S. Patent Nos. 9,369,741 (IPR2019-00231), 9,621,956 (IPR2019-00281), and 9,118,948 (IPR2019-01421) [Comcast Cable Communications, LLC v. Rovi Guides Inc.]. Finished 07/20.
- Depositions and claim construction tutorial in Netfuel, Inc. v. Cisco Systems, Inc. (5:18-cv-2352-EJD, N.D. Cal). Finished 08/20.
- A deposition in Inter Partes Review of U.S. Patent No. 9,663,659 (IPR2019-00992) [Cisco Systems, Inc. v. NetFuel, Inc.]. Finished 09/20.
- A deposition in Monarch Network Solutions, LLC v. Cisco Systems, Inc. (2:20-CV-00015-JRG, E.D. Tex.). Finished 09/20.
- Certain Audio Players and Controllers, Components, Thereof and Products Containing the Same (ITC Inv. No. 337-TA-1191) [Sonos, Inc. v. Google LLC and Alphabet, Inc.]. Finished 02/21.
- A deposition and trial testimony in Gigamon, Inc. v. Apcon, Inc. (2:19-cv-300-JRG, E.D. Tex.). Finished 04/21.
- Depositions and trial testimony in TQ Delta, LLC v. 2Wire, Inc. (1:13-cv-01835-RGA, D. Del.).

- A deposition in Finjan, Inc. v. Cisco Systems, Inc. (5:17-cv-00072-BLF-SVK, N.D. Cal).
- Depositions in TQ Delta, LLC v. AdTran, Inc. (14-cv-954-RGA, 15-cv-121-RGA, D. Del).
- A deposition in Sprint Communications Company LP v. Charter Communications, Inc. et al. (17-1734-RGA, D. Del.).
- A deposition in Contour IP Holding, LLC v. GoPro, Inc. (17-CV-04738-WHO, N. D. Cal.)
- A deposition in Finjan, Inc. v. Rapid7, Inc. (1:18-cv-01519-MN, D. Del.)
- A deposition in Motorola Solutions, Inc. v. Hytera Communications Corp. LTD (1:17-cv-01972, N.D. Ill.)
- A deposition in Finjan, Inc. v. SonicWall, Inc. (5:17-cv-04467-BLF-HRL, N.D. Cal.)
- A deposition in Shopify, Inc. v. Express Mobile, Inc. (1:19-cv-00439-RGA, D. Del.)
- A deposition in Icon Health and Fitness, Inc. v. Peloton Interactive, Inc. (20-1386-RGA, D. Del.)
- A deposition in Linksmart Wireless Technology, LLC v. Gogo, LLC and Panasonic Avionics Corp (8:18-CV-00654-JAK-JDE and 8:18-CV-00662-JAK-JDE)
- Depositions in Luminati Networks Ltd. v. Teso LT, UAB, Oxysales, UAB, and Metacluster LT, UAB, and Luminati Networks Ltd. v. Code 200, UAB, Oxysales, UAB, and Metacluster LIT, UAB (2:19-cv-00395-JRG, E.D. Tex and 2:19-cv-00396-JRG, E.D. Tex.)

Appendix B

PART

1

Packet Network Foundations

Part I of this book introduces four widely deployed packet network technologies: X.25, frame relay, asynchronous transfer mode (ATM), and Internet protocol (IP).

Before packet networks, communications technology used circuit-switched telephone networks with dedicated, analog circuits that functioned on a “always on once activated” basis. A dedicated circuit cannot be used for other purposes even if no communications are taking place at the moment. In regard to telephone conversations, it is estimated that on the average a dedicated circuit carried active traffic only 20 to 25 percent of the time and is idle the other 75 to 80 percent. Moreover, other services such as video data streams cannot be efficiently carried on circuit-switched networks.

Packet networks based on packet switching technologies represent a radical departure. The key idea behind packet switching is that a message or a conversation is broken into independent, small pieces of information called *packets* that are either equal or variable in size. These packets are sent individually to a destination and are reassembled there. No physical resource is dedicated to a connection, and connections become virtual, thus allowing many users to share the same physical network resource.

The concept of packet switching is attributed to Paul Baran who first outlined its principles in an essay published in 1964 in the journal *On Distributed Communications*. The term *packet switching* itself was coined by Donald Davies, a physicist at the British National Physical Lab, who came up with the same packet switching idea independently. It is interesting to note that a few decades earlier, a similar discovery in physics by Albert Einstein—that waves of light can be broken into a stream of individual photons—led to the development of quantum mechanics.

Packet networks allow more efficient use of network resources. Each packet occupies a transmission facility only for the duration of the transmission, leaving the facility available for other users when no transmission is taking place.

Packet-switched networks are highly fault-tolerant. From the very start of their development, network survivability was a major design goal. Because packet networks do not rely on dedicated physical connections, packets can be routed via alternative routes in case of an outage in the original communications link.

Packet networks can support bandwidth on-demand and flexible bandwidth allocation. Bandwidth is allocated at the time of communication, and the amount of bandwidth allocated is based on need. In

CHAPTER
8

Local Area Networks

8.1 Introduction

A local area network is a high-speed data network that covers a relatively small geographic area. It typically connects workstations, personal computers, printers, servers, and other end-user devices, which are collectively also known as *data terminal equipment*. The common applications of LAN include shared access to devices and applications, file exchange between connected users, and communication between users via electronic mail and others. LANs are also private data networks, because they belong to an organization and are used to carry data traffic as opposed to voice traffic.

This section provides a brief introduction to LAN history, standards, protocol stacks, topologies, and devices.

8.1.1 LAN History and Standards

LAN is a type of broadband packet access network that carries the packet data traffic of an organization. LAN interconnects the end users of an organization to an outside public data network such as the Internet.

The basis of LAN technologies and standards was defined in the late 1970s and early 1980s. LAN technologies really emerged with the Internet itself, and the first widely deployed LAN technology, Ethernet, is almost as old as the Internet itself. The overwhelming majority of the deployed LANs are Ethernet.

IEEE 802, a branch of the International Institute of Electrical and Electronics Engineers (IEEE), is responsible for most of the LAN standards. These standards have also been adopted by other standards organization such as ANSI and ISO. The major LAN standards are listed in Table 8-1.

8.1.2 LAN Protocol Stacks

The LAN protocols operate at the bottom two layers of the OSI network reference model, i.e., at the physical layer and the data link layer, as shown in Fig. 8-1. The physical layer is primarily concerned with the transmission medium and its physical characteristics for digital signal transmission. The data link layer consists of two sublayers, the medium access control (MAC) sublayer and logical link control (LLC) layer. The MAC sublayer is responsible for controlling access to a shared medium by multiple users simultaneously. The LLC sublayer is responsible for

TABLE 8-1

IEEE 802 LAN Standards Summary

IEEE 802 specification	LAN technology	Description
IEEE 802.1 (ISO 15802-2)	General information	Details how the other 802 standards relate to one another and to the ISO OSI reference model.
IEEE 802.2 (ISO 8802.2)	LLC framework	Divides the OSI data link layer into two sublayers and defines the functions of the LLC and MAC sublayers
IEEE 802.3	Ethernet	Defines the CSMA/CD protocol, which is used in Ethernet applications and has become synonymous with Ethernet
IEEE 802.4 (ISO 8802.4)	Token bus	Defines the token-passing bus access method
IEEE 802.5 (ISO 8802.5)	Token ring	Defines the Token Ring access method
IEEE 802.7	Broadband LAN	Recommended practices for broadband LANs
IEEE 802.11	Wireless LAN	Wireless LAN medium access control (MAC) and physical layer specifications
IEEE 802.15	Wireless personal area network (WPAN)	WPAN MAC and physical layer specifications
IEEE 802.16	Broadband fixed wireless metropolitan area networks (MANs)	Air interface specification for fixed broadband wireless access systems
IEEE 802.12	100 VG-AnyLAN	Defines a LAN technology that supports the operations of any existing LAN protocol, including the Ethernet frame format and Token Ring frame format, but not both at the same time

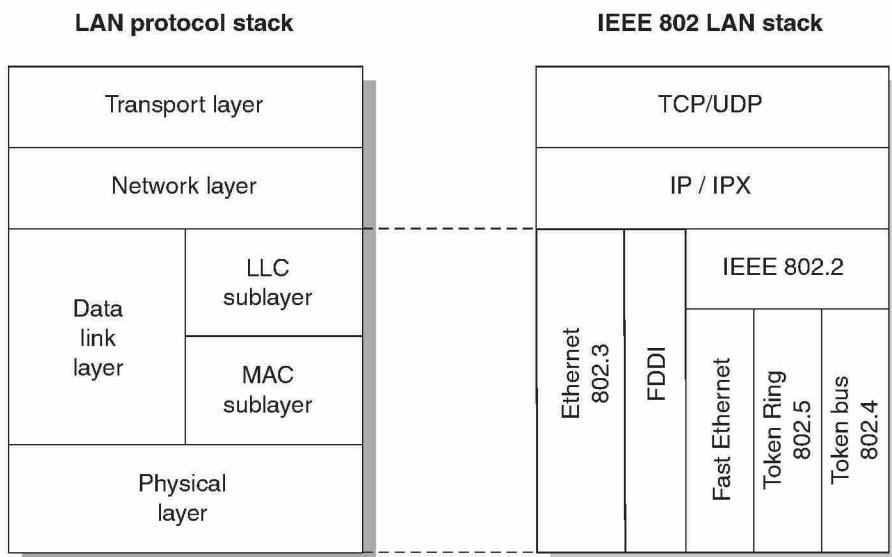
interfacing to the upper layers, such as IP and the Internetwork Packet Exchange protocol (IPX). Any layer above the data link layer is beyond the scope of the LAN protocols.

The IEEE 802 LAN standards are compatible at the upper part of the data link layer, i.e., at the LLC sublayer, but differ from each other at the MAC sublayer and physical layer.

The scope of each LAN protocol may vary. Some cover the entire two bottom layers. For example, Ethernet as defined in IEEE 802.3 (IEEE 2001) and FDDI as defined in IEEE 802.5j (IEEE 1998c) cover the physical layer and both sublayers of the data link layer, as shown on the right-hand side of Fig. 8-1. Other LAN protocols, such as Token Ring and token bus,

Figure 8-1

The LAN protocol stack.



specify the physical layer and the MAC sublayer while sharing a common LLC specification defined in IEEE 802.2, as shown on the right-hand side of Fig. 8-1.

8.1.2.1 Physical Transmission Medium The LAN transmission medium can be divided into the two general categories of wired and wireless. This chapter focuses only on the wired or wireline LAN technology, while Chap. 9 will describe wireless LAN.

There are basically three types of transmission media used in wireline LAN deployment: copper twisted pair, coaxial cable, and optical fiber. The type of transmission medium determines the data rate and transmission distance.

TWISTED PAIR COPPER WIRE Twisted pair, both shielded and unshielded, is a pair of copper wires that are twisted to increase the transmission distance. It is the least costly of the three wireline LAN media, and one of the most common transmission media currently used in LAN applications. It is primarily used in star and hub LAN configurations in office buildings. The maximum transmission distance of twisted pair cable depends on the target data rate; typically the limit is 100 m without repeater. The data rate of copper twisted pair normally is not as high as that of other transmission media and depends on factors such as transmission distance and the modulation scheme used for transmission. The

longer the transmission distance is, the lower the bit rate is. It is not uncommon to see twisted pair achieve a bit rate of over 1 Mbps for a distance of 100 meters.

COAXIAL CABLE Coaxial cable, whose transmission wire is insulated with dielectric insulating material and braided out conductor, can achieve higher data rates and longer transmission distances. There are two kinds of coaxial cables: thin wire and thick wire, referring to the difference in the cable diameters, thin wire being 0.25 in diameter and thick wire being 0.5 in diameter. Thin-wire coaxial cable reaches shorter distances, typically 200 m with the data rate of over 10 Mbps, while thick-wire cable can reach over 500 m with the same data rate.

OPTICAL FIBER Optical fiber carries data in the form of flashing light beams in a glass fiber, as opposed to electrical signals on a wire. Optical fiber can achieve much higher data rates than coaxial cable or twisted pair over much longer distances. The fiber transmission equipment consists of fiber cable, special electrical-to-optical and optical-to-electrical converters, light emitters such as light-emitting diodes (LEDs) or laser and optical receivers. These transmission components have been much more costly than twisted pair and coaxial cable. However, with the advent of new optical transmission technologies and a massive market for broadband applications, the cost has come down considerably in recent years and the optical fiber is becoming a common choice for LAN deployment.

LANs can use one type of transmission medium or a mix of types. For example, lower-speed twisted pair can be used between a computer and a hub, while coaxial cable can be used between a branch hub and a main hub and high-speed optical fiber cable can be used between a main hub and an outside router.

8.1.2.2 Media Access Control Sublayer A LAN technology must address the issue of resource contention because multiple users share the same transmission medium. A contention occurs when two DTEs transmit data at the same time. There are basically two MAC mechanisms for LAN: carrier-sense multiple access with collision detection and control token.

CSMA/CD The CSMA/CD access control method is used in Ethernet and can be characterized as "listen and send." A network device first listens to the wire when it has data to send, then sends the data when it finds that no other device is sending the data. After it finishes sending the

data, it listens to the wire again to detect if any collision occurs while it transmitted data. A collision occurs when two devices send data simultaneously. If a collision is detected, the device waits for a random amount of time before resending the data. The randomness of the wait period makes the possibility of another collision very small. However, this algorithm is not deterministic, and when the number of users increases to a large enough point, network performance deteriorates drastically owing to the large number of collisions.

The major advantage of CSMA/CD is its simplicity. It is easy to implement and works well in the LAN environment.

CONTROL TOKEN Control token is a special network packet used to control access to a shared transmission medium. A token is passed around a network from device to device. When a device has data to send, it must wait until it has the token, at which time it sends its data. When the transmission is complete, the token is released so that other devices may use the network to transmit their data. A major advantage of token-passing networks is that they are deterministic. In other words, it is easy to calculate the maximum time that will pass before a device has the opportunity to send data. This explains the popularity of token-passing networks in some real-time environments such as factories, where machinery must be capable of communicating at determinable intervals. Token-passing networks include Token Ring and FDDI.

MAC ADDRESS The MAC address is a number that is hard-wired into each LAN card such as the Ethernet Network Interface Card or adapter that uniquely identifies this device on a LAN. The MAC addresses are 6 bytes in length, and are usually written in hexadecimal such as 12:34:56:78:90:AB. The colons in the address may be omitted, but generally make the address more readable. Each manufacturer of LAN devices has a certain range of MAC addresses, just like a range of telephone numbers, that they can use. The first 3 bytes of the address denote the manufacturer.

8.1.2.3 Link Layer Control Sublayer The LLC sublayer, as defined in the IEEE 802.2 standard, mainly hides the differences between various MAC sublayer implementations such as Ethernet, Token Ring, and FDDI and presents a uniform interface to the network layer. This allows different types of LANs to communicate with each other.

The IEEE 802 LLC protocol defines a generic LLC protocol data unit that includes both user data and LLC header. The LLC header contains a control field that in turn contains the fields such as protocol ID and

header type. Also found in the LLC header are source and destination address fields.

8.1.3 Data Transmission Methods

There are three data transmission modes in LAN environments: point-to-point, multicast, and broadcast. In each transmission mode, a single packet is sent to one or more nodes.

In *point-to-point transmission*, which is also known as *unicast*, a single packet is sent from a source to a destination on a LAN. First, the source node addresses the packet by using the address of the destination node. The packet is then sent onto the LAN, and the LAN then passes the packet to its destination.

In *multicast transmission*, a single data packet is copied and sent to a specific subset of nodes on a LAN. First, the source node addresses the packet by using a special type of address, called a *multicast address*. The packet is then sent onto the LAN, which makes copies of the packet and sends a copy to each node that is part of the multicast address.

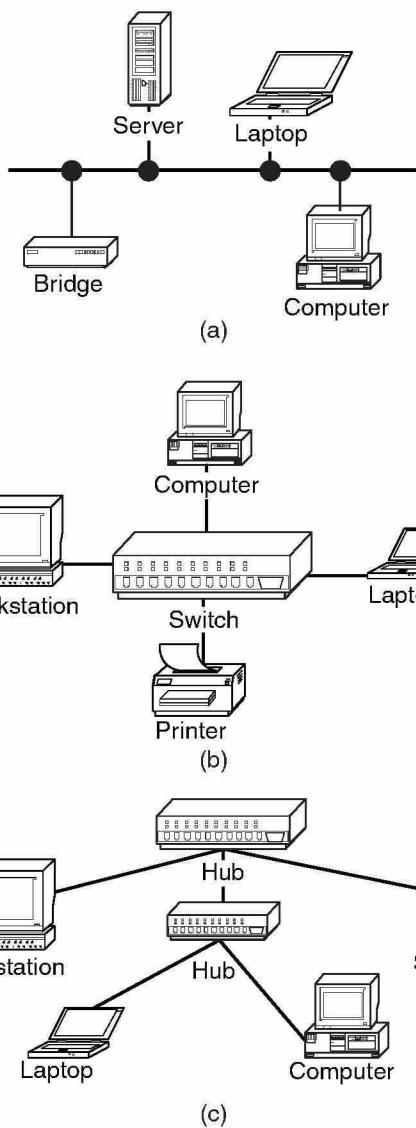
In *broadcast transmission*, a single piece of data is copied and sent to all the nodes on a LAN. In this type of transmission mode, a source node addresses a packet by using a broadcast address. The packet is then sent onto the LAN, which makes copies of the packet and sends a copy to every node on the LAN.

8.1.4 LAN Topology

A LAN topology defines how the data terminal equipment such as desk top computers, printers, and server computers, and LAN internetworking devices such as switches, routers, and hubs, are interconnected to each other. In general, there are four types of LAN topologies: bus, star, ring, and hub. Each has some advantages, which will now be discussed (Halsall 1996).

8.1.4.1 Bus Topology The bus is one of the most common LAN topologies. A simple bus topology is characterized by a central cable that runs through end-user equipment like computers and servers, as shown in Fig. 8-2(a). A physical connection, also known as a *tap*, is made to the cable for each user terminal or computer to access the network. MAC circuitry and the software implementing the control scheme together allow the connected users to share the common cable and transmission

Figure 8-2
Bus, star, and hub
LAN topology
examples.



bandwidth. A slightly more complicated bus topology may consist of multiple layers of buses. A bus cable can be connected to another bus cable, which in turn may be connected to yet another cable. This forms a topology that looks like an uprooted tree.

8.1.4.2 Ring Topology Ring-based LAN topology is characterized by a cable that passes from one DTE to another until all the DTEs are connected to form a ring or loop. A distinct feature of ring topology is

that traffic travels in one direction only. Between two neighboring user DTEs, it is a direct point-to-point link that carries traffic in one direction only, termed *unidirectional*. Again, medium access circuitry and a control algorithm are built into a DTE and network to allow each DTE a fair chance to access the cable ring.

8.1.4.3 Star Topology In star topology, there is a focal point that is either a switch or a router, and all end-user DTEs are connected to the central point via a point-to-point cable, as shown in Fig. 8-2(b). This is a typical voice-service PBX configuration that is also used to interconnect end-user DTEs, although not as common as other topologies. Compared to the other topologies, star topology has more complicated wiring.

8.1.4.4 Hub Topology A fourth common topology is the hub structure, which is a mix of the ring and bus topologies. A hub topology is simply a bus or ring wiring collapsed into a central unit. A hub does not perform any switching or intelligent processing. All a hub does is simply retransmit all signals received from a DTE to all other DTEs with a set of repeaters inside the hub. As shown in Figure 8-2(c), a hub can be connected to another hub to form a hierarchy of hubs and DTEs that looks like a tree structure.

8.1.5 LAN Internetworking Devices

An internetworking device interconnects two or more other LAN devices. Based on functionality, there are three types of such internetworking devices: repeater, bridge, and router or switch.

8.1.5.1 Repeater A LAN *repeater* is a physical layer device used to connect two LAN cable segments so a LAN will extend further in distance. A repeater essentially boots digital signals to allow a series of cable segments to be treated as a single cable. It receives signals from one network segment and amplifies, retimes, and retransmits those signals to another network segment. A LAN repeater operates at the physical layer without any intelligence to perform any filtering or other traffic processing. In addition, all electrical signals, including electrical noise and errors, are repeated and amplified as well. The total number of repeaters within a LAN is limited due to timing and other issues.

8.1.5.2 LAN Hub A *hub* is a physical layer device that connects multiple user stations, each through a dedicated cable. In some respects, a

hub functions as a multiport repeater. Hubs are used to create a physical star network while maintaining the logical bus or ring configuration of a LAN.

8.1.5.3 LAN Bridge A LAN *bridge* is an internetworking device that interconnects two LAN segments at the data link layer as opposed to the physical layer in the case of a repeater. A bridge must have at least two ports, one receiving incoming frames and one sending outgoing frames. A bridge uses a MAC address to route frames from one segment to another, or even to a different LAN that is the same or different at the physical or MAC layer.

8.1.5.4 LAN Router and Switch A LAN *router* operates at the network layer, interconnecting like and unlike devices attached to one or more LANs. LAN routers normally also support link layer bridging in addition to network layer routing.

A LAN router, as described in Chap. 4 on IP networks and Appendix A, employs routing protocols to dynamically obtain knowledge of destination address prefixes across an entire set of internetworked LANs. A LAN router normally has a packet-forwarding engine that uses a lookup table to identify the physical interface of the next hop toward the destination.

8.2 Ethernet

Ethernet is almost as old as the Internet itself. Since its inception at a Xerox lab in the early 1970s, it has been the dominant protocol for local area networks. By various estimates, Ethernet accounts for somewhere between 80 to 95 percent of worldwide LAN installations.

This section, after first providing some background information, introduces three generations of Ethernet: 10Base Ethernet, Fast Ethernet, and optical Ethernet, with an emphasis on the first two. Gigabit Ethernet and 10 Gigabit Ethernet were described in detail in Chap. 7 in the context of optical transport network, and will be discussed briefly in this chapter in the context of LAN technology.

What is remarkable about Ethernet is its continuity and simplicity. The fundamentals of Ethernet such as Ethernet frame and logical link control, which were already defined for the first generation of Ethernet, have remained largely intact through the rapid technological evolution

of the past two decades. Ethernet is viewed as a kind of plug-and-play technology because it is relatively simple and can operate with very little manual intervention for configuration and provisioning.

8.2.1 Ethernet Basics

8.2.1.1 A Brief History Ethernet was originally developed by Digital, Intel, and Xerox (DIX) in 1972 and was designed as a “broadcast” system where stations on a network can send messages at will. All the stations may receive the messages, but only one specific station to which the message is directed will respond. Robert Metcalf and David Boggs of Xerox are credited with coming up with first Ethernet design. Ethernet was originally designed to run on any medium (copper wire, fiber, or even radio wave), which is where *Ether* in the term *Ethernet* comes from.

The original version of Ethernet was adopted by IEEE Committee 802.3 (IEEE Project 802 was named after the time Ethernet was set up, in February 1980), and the packet format was standardized, which is known as the IEEE 802.3 Ethernet frame.

The Ethernet evolution, based on the transmission technologies and speed, involved at various times in its development, can be divided into the following periods:

- 10BaseT Ethernet, starting from 1972 to the mid-1990s
- 100Base Ethernet, starting from the mid-1990s
- 1000Base Ethernet, starting from 1998
- 10Gig Ethernet, starting from 2000

An Ethernet version is represented in terms of the transmission speed, the transmission medium, and maybe the transmission distance. The prefix number in an Ethernet version such as 10 in 10Base or 100 in 100Base refers to the transmission speed of 10 Mbps and 100 Mbps. The suffix letter refers to the medium type, while suffix number for earlier versions of Ethernet refers to the maximum transmission distance. For example, the letter T in 10BaseT refers to “twisted pair” copper wire and the number 5 in 10Base5 refers to the transmission distance in hundreds of meters.

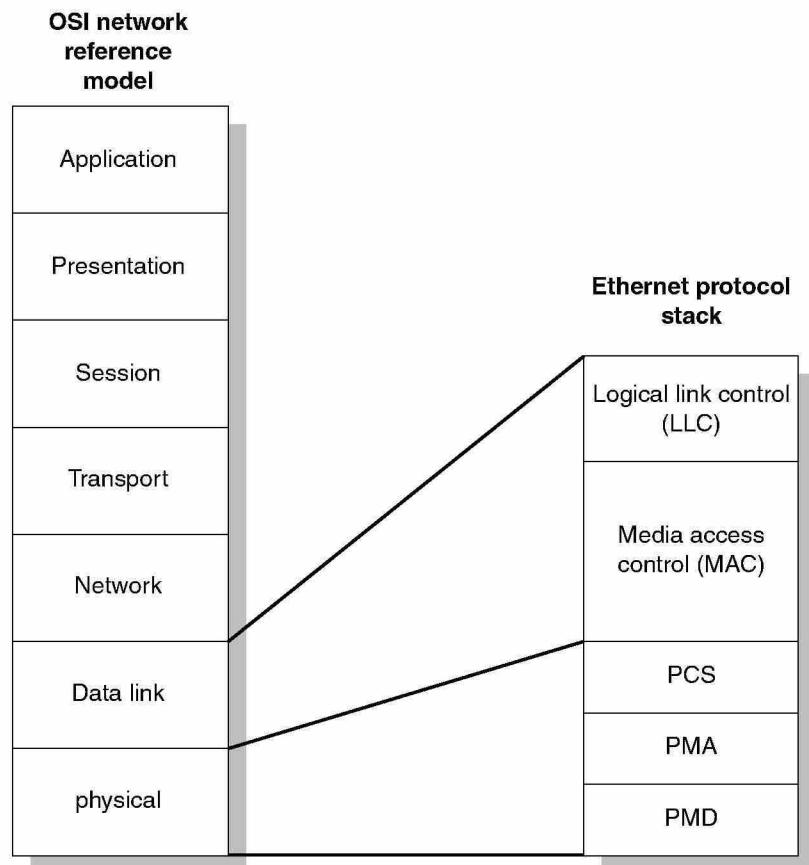
8.2.1.2 Ethernet Protocol Stack The Ethernet protocol stack is similar to the general LAN protocol stack as described earlier. It covers the layers 1 and 2 of the OSI network reference model. In addition, Ethernet further

defines three sublayers for the physical (PHY) layer: PMD, PMA, and PCS, which are briefly discussed here (IEEE 2001c).

The physical medium-dependent (PMD) sublayer defines the Ethernet cables, wiring, and other transmission medium-related components. The physical medium attachment (PMA) defines the type of connectors used to connect an Ethernet device such as an Ethernet NIC, hub, or switch to the Ethernet cable. The physical coding sublayer (PCS) defines a scheme appropriate to the medium to encode and decode data received from/sent to the PMD sublayer (Spurgen 2000).

The Ethernet data link layer, like that of other LAN technologies, is broken into two sublayers: the LLC on the upper half and the MAC on the lower half. The MAC deals with getting data on and off the wire and media access control, as shown in Fig. 8-3. The logical link control

Figure 8-3
The Ethernet protocol stack in reference to the OSI network reference model.



on the upper half of the data link layer deals with error checking and providing a uniform interface to the network layer above.

8.2.1.3 Ethernet Operation Mode Ethernet supports either half-duplex, full-duplex, or both operation modes. Early Ethernet supports only the half-duplex mode of operation, where a station can transmit or receive data but not at the same time. In contrast, a station supporting the full-duplex mode of operation can transmit and receive data simultaneously. It was with the development of Fast Ethernet that Ethernet became able to support both half-duplex and full-duplex modes of operation.

8.2.2 First Generation—10BaseT Ethernet

10BaseT is one of the most popular versions of the first generation of Ethernet, and defines the fundamentals of Ethernet technology upon which later generations of Ethernet have been built. This discussion will cover the area of the physical layer, the media access control sublayer, and the logical link control sublayer.

8.2.2.1 Physical Layer of Ethernet The characteristics of the first generation of Ethernet are summarized in Table 8-2, which includes the transmission medium, transmission distance and data rate, and operation mode.

TABLE 8-2
10Base Ethernet
Summary

Standards	IEEE standard—year first released	PMD type	Date rate	Maximal distance in meters	
				Half duplex	Full duplex
10Base5	802.3—1983	Coax cable (thick Ethernet)	10 Mbps	500	Not supported yet
10Base2	802.3—1985	Coaxial cable (thin Ethernet)	10 Mbps	185	Not supported yet
10Base5	802.3—1987	2 pairs of twisted telephone cable	1 Mbps	250	Not supported yet
10Base-T	802.3—1990	2 pairs of category 3 or better UTP cable	10 Mbps	100	100
10Base-FL	802.3—1993	Two optical fibers	10 Mbps	2000	>2000

SOURCE: IEEE 2000.

The transmission medium has evolved from the original thick coax (10base5) to twisted pair copper wire and then to fiber. Twisted pair is the most common choice of cable for the first generation of Ethernet. Unshielded twisted pair (UTP) is one kind of twisted pair that has two copper wires twisted together and is relatively immune to noise.

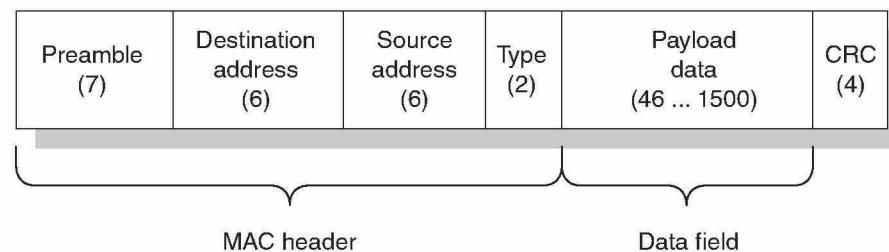
The physical coding sublayer uses Manchester coding, a common coding scheme at the time of first-generation Ethernet that divides each bit into two halves. A 1 is defined by a transition from “low” to “high” in the middle of the bit period, and a 0 is defined as a transition from “high” to “low” in the middle of the bit period.

Most versions of first-generation Ethernet support only the half-duplex mode of operation and have a transmission distance of around 250 m.

8.2.2.2 Ethernet Frame The Ethernet frame defines a structure to hold user data and to be carried on the physical medium. It consists of two parts: a header and the payload data. Figure 8-4 shows the IEEE 802.3 Ethernet frame format, which includes the following:

- *Preamble*. A 7-byte field containing a series of alternating 1s and 0s used by an Ethernet receiver to acquire bit synchronization and frame timing information. This field is generated by the hardware in an Ethernet device.
- *Destination address*. The MAC address of a receiving Ethernet device.
- *Source address*. The MAC address of a sending device.
- *Type*. A 2-byte field indicating the type of data encapsulated, e.g., IP, ARP, RARP, etc.
- *Payload data*. The data field with length ranging from 46 to a maximum of 1500 bytes.
- *Cyclical redundancy check (CRC)*. A 4-byte field used for error detection.

Figure 8-4
The IEEE 802 Ethernet frame structure.



8.2.2.3 Media Access Control Ethernet MAC uses CSMA/CD for access control. By means of carrier sense multiple access, with collision detection, an Ethernet device does the following:

1. Listens to the line before putting a packet “on the wire,” and if the line is busy, waits for a predetermined number of seconds before retry
2. When the line becomes idle, transmits while monitoring for collisions
3. If a collision is detected, sends the jam signal and waits for an algorithmically determined number of seconds before resending any packets
4. If the maximum number of transmission attempts is reached, gives up

8.2.3 Second Generation—Fast Ethernet

Fast Ethernet was defined to meet the demands of fast-growing Internet traffic. In the face of fast growth, 10 Base Ethernet became too slow by the early 1990s to meet all the needs of the Internet’s data traffic flow. The IEEE reconvened the IEEE 802.3 committee in 1992 to upgrade Ethernet to 100 Mbps. Two competing proposals emerged in the process: one simply aimed at increasing the speed of the existing Ethernet as defined by IEEE 802.3 to 100 Mbps, while the other reworked the old Ethernet with a new architecture. The first proposal resulted in the updated IEEE 802.3 specifications, also known as *Fast Ethernet*, that were approved in 1996. The second resulted in the establishment of the IEEE 802.12 committee and the 802.12 standard specifications in 1995, also known as *100VG-AnyLAN*. This subsection briefly describes Fast Ethernet, while the following subsection discusses 100VG-AnyLAN.

One major change in the Fast Ethernet specifications is that shared medium topologies like the bus topology are eliminated in favor of the star topology in order to decrease transmission collisions and increase network throughput. At the center of the star topology is a switching hub that supports full-duplex operation.

8.2.3.1 Physical Layer The Fast Ethernet specifications define three physical media, or physical medium-dependents: 100Base-T4, 100BaseSE-TX, and 100Base-FX. The 100Base-T4 uses four unshielded

twisted pairs of cable to connect a user station to a hub, a very common situation in office buildings. The 100Base-TX uses two pairs of category 5 unshielded twisted pairs. The 100Base-FX uses a pair of optical fiber cables that are defined by ANSI standards for FDDI. Table 8-3 summarizes the Fast Ethernet physical layer characteristics.

Fast Ethernet adopts a faster coding scheme at the physical signaling sublayer, i.e., the 4-bit/5-bit scheme that uses groups of four data bits as a transmission unit, also called an *encoded symbol*, with the fifth bit as the delimiter.

8.2.3.2 MAC Layer Fast Ethernet retains the original Ethernet MAC layer. All the original frame formats, procedures, and media access control algorithms, i.e., CSMA/CD, remain almost identical. This enables the first-generation of 10-Mbps Ethernet LANs to run over 100 Mbps Fast Ethernet without any changes.

8.2.4 100VG-AnyLAN

The IEEE 802.12 standards, originally approved in 1995, were the result of a competing proposal for upgrading the first generation of Ethernet. The central idea behind 100 VG-AnyLAN is to define a LAN technology that supports the operations of any existing LAN protocol, be it Ethernet frame format and Token Ring frame format, but not both at the same time. The main goals of 100VG-AnyLAN include avoiding the frame collisions of the traditional CSMA/CD access method and providing

TABLE 8-3

100Base Ethernet Summary

Standards	IEEE standard—year first released	PMD type	Maximal distance in meters	
			Half duplex	Full duplex
100Base-TX	802.3—1995	Two pairs of category 5 UTP cable	100	100
100Base-FX	802.3—1995	Two optical fibers	400	2000
100Base-T4	802.3—1995	Four pairs of category 3 or better UTP cable	100	Not supported
100Base-T2	802.3—1997	Two pairs of category 3 or better UTP cable	100	100

SOURCE: IEEE 2000b.

prioritized services on LAN (IEEE 1998a). 100VG-AnyLAN did not achieve wide acceptance in the market, largely due to the overwhelming dominance of Ethernet.

The prioritized service is implemented via a demand priority protocol that utilizes a round robin polling scheme for each station to request a priority for each MAC frame from the repeater. Higher priority is given to delay-sensitive frames, while the best-effort service is given to the rest of the frames.

Collision avoidance is achieved via the exclusive use of a switching hub as opposed to the shared media used by traditional Ethernet. A station can transmit only after it is granted permission to do so by the connected repeater. Thus the access control method is deterministic with no collisions.

8.2.5 Gigabit and 10 Gigabit Ethernet

Gigabit Ethernet and 10 Gigabit Ethernet as transport technologies are introduced in Chap. 7 on optical ethernet, which focuses on the physical layer of both Ethernet technologies. This subsection provides an overview of Gigabit Ethernet and 10 Gigabit Ethernet from the perspective of LAN.

8.2.5.1 Gigabit Ethernet Soon after the Fast Ethernet standards were finalized, the work on 1000Base Ethernet began at the IEEE 802.3z committee. After the specifications were completed, large-scale deployment soon followed.

One primary goal of Gigabit Ethernet, like its predecessor Fast Ethernet, is to alleviate the bandwidth crunch on LANs with 10-fold increase in speed. Gigabit Ethernet also preserves the standard 802.3 Ethernet frame format and the minimum and maximum sizes of the frame, so that it is backward-compatible with 100BaseT and 10BaseT Ethernet.

Gigabit Ethernet supports both full- and half-duplex operations, the same as Fast Ethernet. For half-duplex operations, CSMA/CD is used. For full-duplex operations, the standard flow control defined in IEEE 802.3 is used (IEEE 2001b). At the physical layer, Gigabit Ethernet supports both fiber and copper wire as physical media, although optical fiber is the common choice. It uses the recently defined ANSI Fibre Channel standards as the basis for fiber-based media (ANSI 1998).

Gigabit Ethernet equipment, like Ethernet switch or router equipment, is mainly used for LAN backbone, interconnecting distributed multiple LANs, or connecting a LAN to a backbone IP network.

8.2.5.2 10 Gigabit Ethernet Efforts on the 10 Gigabit Ethernet specifications by the IEEE 802.1ae committee were initiated soon after the Gigabit Ethernet specifications were completed. The 10 Gigabit technology clearly targets LAN, the traditional space of Ethernet, and the space beyond LAN such as WAN and MAN. 10 Gigabit Ethernet defines two families of physical layer interfaces: one for local area networks operating at a data rate of 10 Gbps, and one for wide area networks, operating at a data rate compatible with the payload rate of OC-192c/SDH VC-4-64c. 10 Gigabit Ethernet preserves the standard 802.3 Ethernet frame format and the minimum and maximum sizes of the frame, so that it is backward-compatible with 100BaseT and 10BaseT Ethernet, like Gigabit Ethernet.

One important feature of 10 Gigabit Ethernet is that it supports full-duplex operation only. The traditional Ethernet half-duplex operation for shared connections and CAMA/CD is abandoned.

8.3 Token Ring LAN

Token Ring LAN technology was originally developed by IBM in the 1970s, was originally standardized by the IEEE as the standard IEEE 802.5 in 1985, and then was adopted as ISO 8802.5 (IEEE 1998c). The IEEE 802.5 specification is almost identical to IBM's Token Ring network, with some minor differences. Throughout this chapter, the term *Token Ring* generally is used to refer to both IBM's Token Ring network and IEEE 802.5 network unless noted otherwise.

The Token Ring network is well suited for use in commercial and industrial environments, where predictability of the performance is expected.

8.3.1 Transmission Medium

IBM Token Ring uses twisted pair copper wire as the transmission medium even though IEEE 802.5 does not specify a media type. In more recent deployments, optical fiber cable is also used to extend the size of the ring interconnecting hubs beyond their normal limitations.

With unshielded twisted pair, a very common wiring choice, a Token Ring network can have a maximum of 72 stations or nodes, although in practice the number of nodes is normally smaller. With shielded twisted pair (STP) wiring, the number of attached stations can increase up to 250

in theory. The typical distance of a Token Ring LAN, called an *average ring length* (ARL), is about 100 m, and this distance can be extended 10-fold if optical fiber cable is used between hubs.

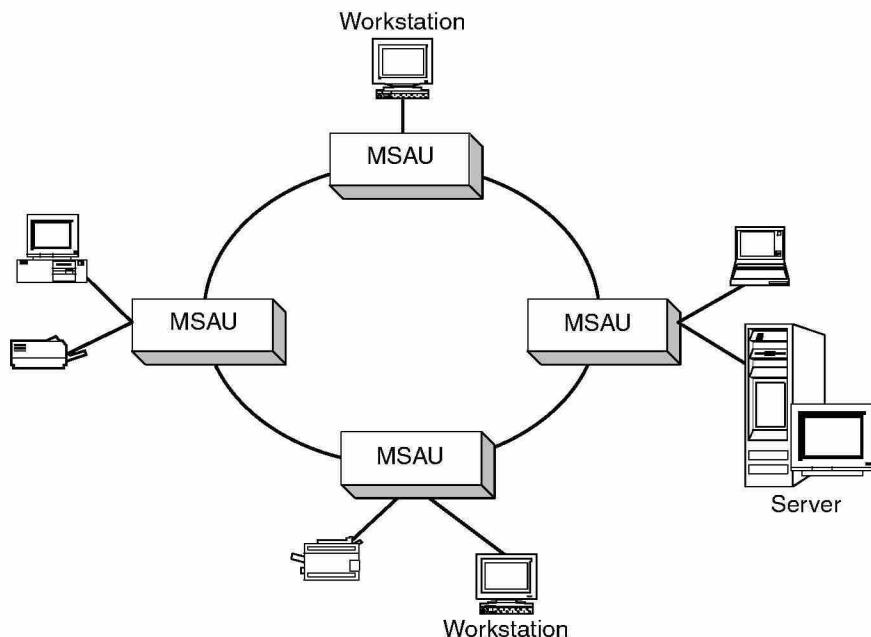
The original IEEE 802.5 Token Ring LAN operates at 4 Mbps, but the standard now covers transmission rates up to 16 Mbps.

8.3.2 Token Ring LAN Configuration and Topology

A Token Ring network typically features a ring topology formed from a set of small clusters or stars, as shown in Fig. 8-5. At the center of each star is a multistation access unit (MSAU) with a set of Token Ring stations connected to it. An MSAU is basically a hub device, and each station is connected to it via a twisted pair cable with two wire pairs. One pair receives data and the other is for transmitting data. The MSAUs are connected together with patch cable or optical fiber cable to form a ring.

An MSAU can be passive or active. A passive MSAU merely provides an electrical path for the data to pass through. An active MSAU amplifies the signals passing through it. With active MSAUs, a Token Ring network can extend further in distance.

Figure 8-5
A Token Ring network topology.



8.3.3 Media Access Control and Frame Format of Token Ring

As the name of the protocol suggests, the media access method used with Token Ring networks is called *token passing*. This is a deterministic access method that ensures no collisions will occur because only one station can transmit at any given time.

There are two types of frame for Token Ring LAN: token frame and data frame. A token frame is a short frame three octets in length, and can turn into a data frame when the token bit is set to 1, as shown in Fig. 8-6.

The token frame has a start delimiter (SD) and an end delimiter (ED), each with a length of one octet. The access control octet has four fields: priority, token indicator, monitor, and reserved bits. The priority field indicates the frame priority and a station can seize the token only if its own priority is equal to or higher than the token priority. The token indicator bit indicates whether the frame is a token or a data frame. The monitor field prevents any frame from circulating on the ring endlessly. The reserved bits field allows a station with higher priority to reserve the next token to be issued with the indicated priority.

A data frame is a superset of the token frame with additional fields such as destination and source addresses, data, and FCS fields, as shown in Fig. 8-6.

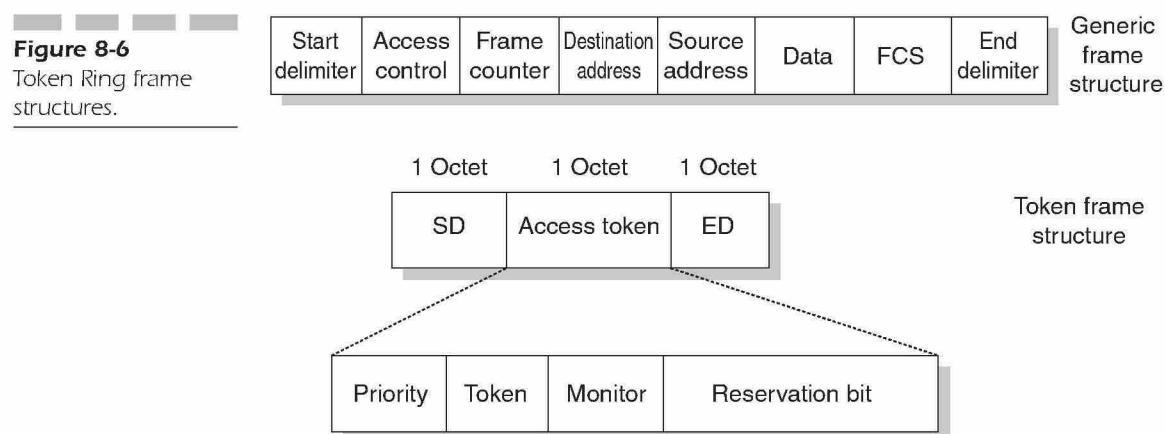
8.3.4 Token Ring LAN Operations

User data travels on the Token Ring network in one direction only, as in other ring topologies. The token frame is passed from node to node. Possession of the token gives a station permission to transmit data.

8.3.4.1 Data Transmission Operations When a station has data to send, it seizes the token first, then changes the token indicator bit to turn the token into the start of a data frame sequence. It then inserts user information and a destination address into the frame and sends it to the next station downstream on the ring.

Each station on the ring examines the data frame and passes it onto the next neighboring station if it is not the intended destination station. The destination station copies the frame for further processing, and sets a bit in the frame to acknowledge receipt of the frame.

The frame continues to circulate the ring until it reaches the sender. The sender removes the frame when it finds that the frame has been “seen” by the destination station. When the sender finishes sending the



last frame, it regenerates the token and puts it on the ring to allow other stations on the network to transmit data.

When a data frame is on the ring, no token frame is on the ring at the same time. This prevents two stations from transmitting data simultaneously so that no collisions occur.

8.3.4.2 Priority Token Ring LAN uses a priority system that allows the operator to assign high priority to some stations that can use the network more frequently than others. Briefly, the priority scheme works as follows: The token frame has two fields that control the priority: the priority field and the reservation field. Only those stations with a priority equal to or higher than the priority level contained in the token frame can seize the token. After the token is seized and changed to a data frame, those stations with a priority level higher than that of the transmitting station can reserve the token for the next round of token passing. When the next token frame is generated, it contains the higher-priority level of the reserving station. Once the reserving station finishes sending data, it is responsible for resetting the token frame's priority level to the original level in order to allow other stations a chance to transmit data.

8.3.4.3 Ring Management A station on a Token Ring LAN plays the role of either an active monitor or a standby monitor station. There is only one active monitor on a ring, and it is chosen during a process called the *claim token process*. The active monitor is responsible for maintaining the master clock, issuing a “neighbor notification,” which is similar to a keep-alive message, detecting lost tokens and frames and purging the ring to get rid of endlessly circulating frames. Any station on the

ring can be the active monitor station if the current active monitor goes down, via the same claim token process.

8.4 FDDI

Fiber-distributed data interface (FDDI) LAN is another incarnation of Token Ring LAN, defined by ANSI (ANSI 1987, 1988) to fill two needs at the time the protocol was adopted. FDDI is intended to fill the need for a large amount of bandwidth on enterprise LANs and the need for reliable and fault-tolerant networks when enterprises start moving critical applications onto their networks. It was adopted by IEEE as IEEE 802.5 (IEEE 1998c), and by ISO. All the specifications are compatible and completely interoperable.

The FDDI standards define the physical layer and the data link layer of the LAN protocol stack. Specifically, they consist of four separate specifications covering the LAN physical layer protocol, PMD, MAC, and station management.

8.4.1 FDDI Basics

FDDI uses two types of optical fibers as primary transmission media: single-mode fiber, which is more expensive but has higher capacity, and multimode fiber, which is relatively inexpensive but has less capacity. The FDDI specification allows for 2 km between stations using multimode fiber and a longer distance with single-mode fiber, and supports a data rate of 100 Mbps.

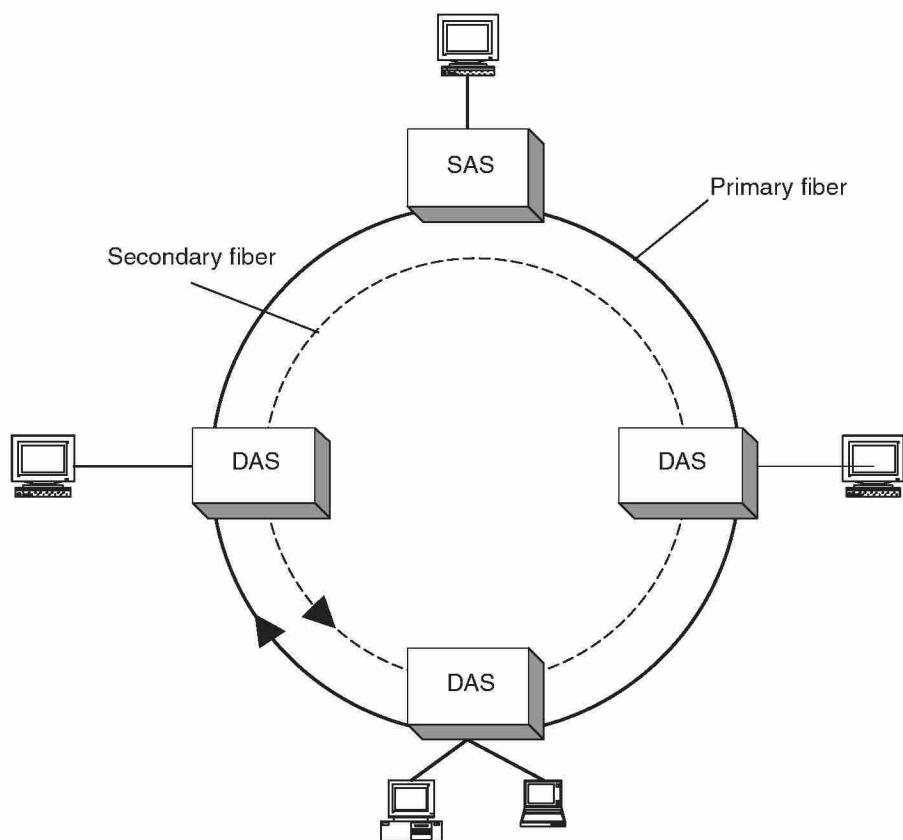
The FDDI frame structure is very similar to that of the Token Ring frame structure described earlier, and it can be as large as 4500 octets. Like the Token Ring frame, the FDDI token frame is a subset of a general frame with three fields: a start delimiter, an end delimiter, and a frame control, which have identical fields to the token frame of Token Ring.

8.4.2 FDDI Configuration and Access Control

FDDI uses two counterrotating rings to enhance its fault tolerance capability: a primary ring and a secondary ring. As shown in Fig. 8-7, the secondary ring can be used to provide additional bandwidth or purely as a backup to the primary ring.

Figure 8-7

Configuration and components of FDDI network.



In an FDDI LAN, there are two kinds of stations: the dual attachment station (DAS), which is connected to both rings, and the single attachment station (SAS), which is attached only to the primary ring. Another FDDI LAN device is the attachment concentrator, which allows multiple DASs or SASs to connect to either ring.

FDDI uses a media access control method that is different from that used by basic Token Ring. As discussed above, Basic Token Ring uses priority and reservation bits in the access control field of the token. In contrast, FDDI uses timed token rotation protocol, which operates as follows: For each rotation of the token, each station computes the time that has expired since it last received the token; this time is called the *token rotation time* (TRT). The TRT includes the time a station needs to transmit any of its waiting frames and the time all other stations in the ring need to transmit any of their waiting frames. TRT will be shorter if the system load is light and longer if the load is heavy. There is a pre-defined parameter called the *target token rotation time* (TTRT). Upon

receipt of a token, a station computes its TRT and the difference between the TTRT and the just computed TRT. The difference, known as the *token hold time* (THT), decides whether and how long the station can transmit the waiting frames. If the THT is positive, the station can spend up to the amount of time equal to the THT in transmitting data. If the THT is negative, the station cannot transmit any frame for this rotation of the token. This time token rotation protocol prevents a station from holding the token for an excessive amount of time and ensures that all stations have a fair chance to use it. This is the same mechanism the token bus protocol uses.

8.4.3 Station Management

There is one management station that acts as the manager on an FDDI ring, and each station has a station management agent. An agent station communicates with the management station to negotiate TTRT and reports the station status.

8.4.4 CDDI

A standard specification similar to FDDI for copper wire has emerged more recently, called the Copper Distributed Data Interface (CDDI) to be consistent with FDDI naming convention. CDDI is an implementation of the FDDI protocol on the copper medium and supports 100 Mbps over a 100-m distance from a desktop to a concentrator (ANSI 1995).

CDDI was defined by the ANSI X3T9.5 committee. It is officially named the Twisted-Pair Physical Medium-Dependent (TP-PMD) Standard to indicate that the focus of the specification is on the twisted pair physical medium, with rest of the protocol including the MAC algorithm and network configurations identical to that of FDDI.

REVIEW QUESTIONS



1. What are the three media types for LAN? Describe the relationships between transmission distance and data rate.
2. The IEEE 802 LAN standards and protocols cover only the bottom two layers of the network reference model. Describe the responsibilities of each of the two bottom layers in the LAN context.

3. Describe the two media access control methods used for LANs and discuss the characteristics of each.
4. Describe the four LAN topologies and explain which ones are most commonly deployed.
5. Describe the operations of a LAN bridge and the differences between a LAN bridge and a LAN router.
6. Describe the responsibilities of the MAC and LLC sublayers in the LAN protocol stack.
7. Explain why Ethernet is a simple technology in terms of access control methods, network topology, and frame format.
8. Describe the differences between Fast Ethernet and the first generation of Ethernet in terms of transmission media, network topologies, and operation modes.
9. Describe the Token Ring LAN topology and how the token is passed around on a Token Ring LAN.
10. Describe the operations of the CSMA/CD access control method and compare it with the token-passing scheme.
11. Describe how the priority scheme in a Token Ring LAN allows some stations to transmit more data than other stations and how to prevent a frame from circulating the ring indefinitely.
12. Describe how the time token rotation protocol works as used in FDDI and token bus networks. Specifically, discuss how it prevents a station from holding onto the token for an excessive amount of time.
13. Briefly describe CDDI and compare it with FDDI.

REFERENCES

ANSI. 1987. "FDDI—Token Ring Media Access Control (MAC)." ANSI X3.139. Web site: www.anci.org.

ANSI. 1988. "FDDI—Token Ring Physical Layer (PHY) Protocol." ANSI X3.148. Web site: www.anci.org.

ANSI. 1995. "FDDI—Token Ring Twisted Pair Physical Layer Medium Dependent (TP-PMD)." ANSI X3.263. Web site: www.anci.org.

ANSI. 1998. "Fibre Channel: Physical and Signaling Interface—3." ANSI INCITS 303. Web site: www.anci.org.

IEEE. 1990. "Token Passing Bus Access Method and Physical Layer Specifications." IEEE 802.4. Web site: www.ieee.org.

IEEE. 1994. "Distributed Queue Dual Bus (DQDB) Access Method and Physical Layer Specifications." IEEE 802.6. Web site: www.ieee.org.

IEEE. 1997. "IEEE Recommended Practices for Broadband Local Area Networks." IEEE 802.7. Web site: www.ieee.org.

IEEE. 1998a. "Demand-Priority Access Method, Physical Layer and Repeater Specifications." IEEE 802.12. Web site: www.ieee.org.

IEEE. 1998b. "Logical Link Control." IEEE 802.2. Web site: www.ieee.org.

IEEE. 1998c. "Token Ring Access Method and Physical Layer Specification." IEEE 802.5. Web site: www.ieee.org.

IEEE. 1999. "Wireless LAN Medium Access Control (MAC) and Physical Layer Specifications." IEEE 802.11. Web site: www.ieee.org.

IEEE. 2001a. "Air Interface for Fixed Broadband Wireless Access Systems." IEEE 802.16. Web site: www.ieee.org.

IEEE. 2001b. "Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications." 2001 edition. IEEE 802.3. Web site: www.ieee.org.

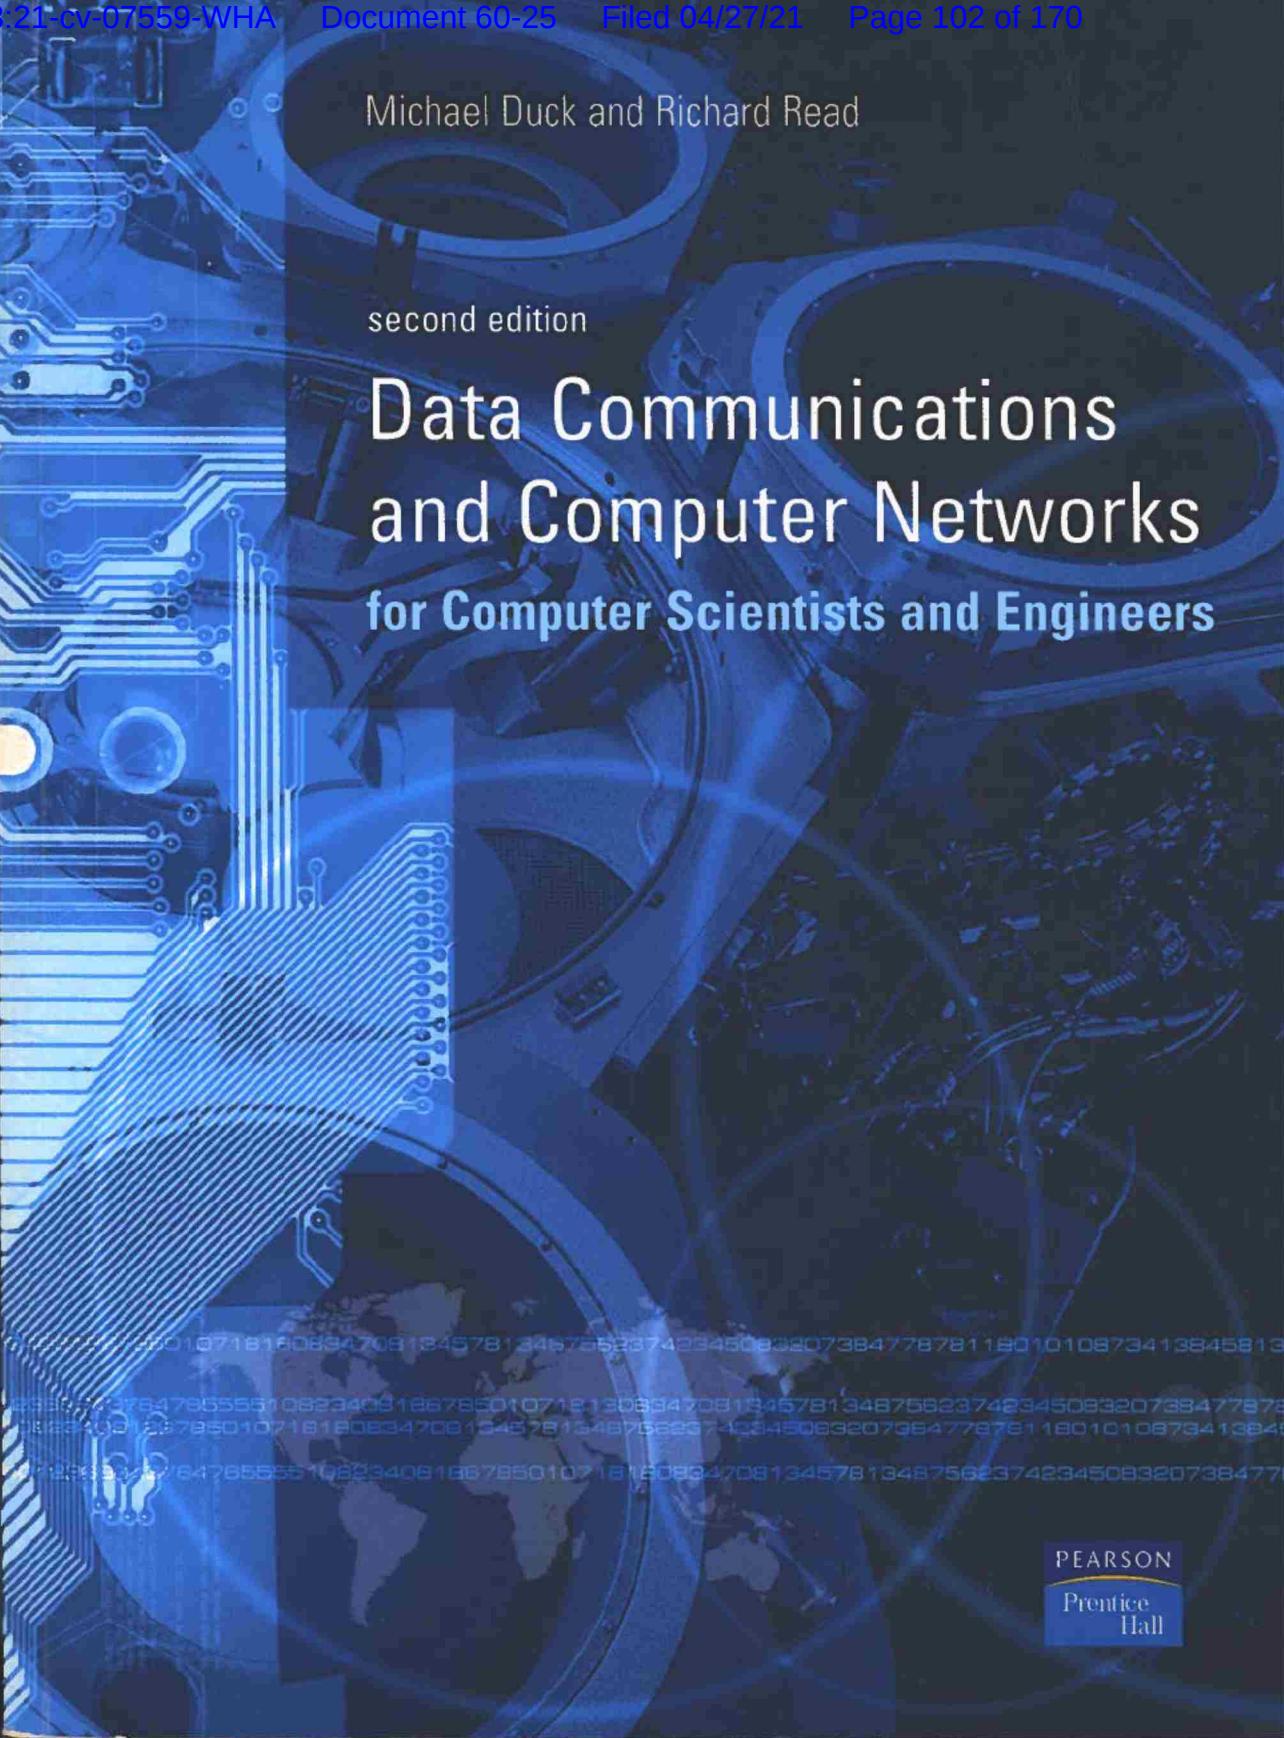
IEEE. 2001c. "Local and Metropolitan Area Networks: Overview and Architecture." IEEE 802. Web site: www.ieee.org.

IEEE. 2002a. "Media Access Control (MAC) Parameters, Physical Layer and Management Parameters for 10Gb/s Operation." IEEE 802.3ae. Web site: www.ieee.org.

IEEE. 2002b. "Wireless Medium Access Control (MAC) and Physical Layer Specifications for Wireless Personal Area Networks (WPANs™)." IEEE 802.15. Web site: www.ieee.org.

Spurgen, C. 2000. *Ethernet: The Definitive Guide*. O'Reilly and Associations.

Appendix C



Michael Duck and Richard Read

second edition

Data Communications and Computer Networks

for Computer Scientists and Engineers

PEARSON
Prentice
Hall

Data Communications and Computer Networks

for Computer Scientists and Engineers

Second Edition

**Michael Duck
and
Richard Read**



Harlow, England • London • New York • Boston • San Francisco • Toronto • Sydney • Singapore • Hong Kong
Tokyo • Seoul • Taipei • New Delhi • Cape Town • Madrid • Mexico City • Amsterdam • Munich • Paris • Milan

Pearson Education Limited
Edinburgh Gate
Harlow
Essex CM20 2JE
England

and Associated Companies throughout the world

Visit us on the World Wide Web at:
www.pearsoneduc.com

First published 1996
Second edition 2003

© Pearson Education Limited 1996, 2003

The rights of Michael Duck and Richard Read to be identified as authors of this work have been asserted by them in accordance with the Copyright, Designs and Patents Act 1988.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without either the prior written permission of the publisher or a licence permitting restricted copying in the United Kingdom issued by the Copyright Licensing Agency Ltd, 90 Tottenham Court Road, London W1T 4LP.

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Pearson Education Limited has made every attempt to supply trademark information about manufacturers and their products mentioned in this book. A list of trademark designations and their owners appears below.

Trademark notice

The following designations are trademarks or registered trademarks of the organizations whose names follow in brackets: Cisco Systems (Cisco Systems, Inc.); Nortel Networks, Passport and Preside (Nortel Networks).

ISBN 0 130 93047 4

British Library Cataloguing-in-Publication Data

A catalogue record for this book is available from the British Library

Library of Congress Cataloging-in-Publication Data

Copy to follow

10 9 8 7 6 5 4 3 2 1
07 06 05 04 03

Typeset in 10/12pt Times by 35

Printed and bound in Great Britain by Henry Ling Ltd, at the Dorset Press, Dorchester, Dorset

**General Library System
University of Wisconsin - Madison
728 State Street
Madison, WI 53706-1494
U.S.A.**

Chapter 9

Introduction to local area networks

9.1	Medium Access Control	200
9.2	LAN performance	206

LANs generally encompass a small physical area, no more than a few kilometres in diameter, and are usually confined within a single site. There may be as many as several hundred stations. LANs use relatively high signalling rates, originally of the order of 10 Mbps, but now typically 100 Mbps, and even as high as 1 Gbps. This chapter is an introduction to LANs and reviews the mechanisms to control station access to the network and also makes some comparisons of their performance. The next chapter will examine the standards for lower speed LAN operation up to a few tens of Mbps. Higher speed LANs, operating at 100 Mbps and beyond, are discussed in Chapter 11.

Messages within a LAN are transmitted as a series of variable length frames using transmission media which introduce only relatively low error rates. Since the length of the medium in a LAN is, unlike in WANs, relatively short, propagation delay tends to be short. Although real-time services such as voice and video are carried over LANs, their frame-based transmissions tend to be more suited to data applications. This is because delay may be unpredictable and hence, on occasion, excessive for real-time use. In addition, transmission capacity has tended to be insufficient to meet the high-data-rate demands of video. However, with the emergence of 100 Mbps operation and beyond, coupled with advances in compression, real-time and multimedia operation is now feasible in many LANs.

A LAN comprises three hardware elements, namely a transmission medium, a controlling mechanism or protocol to govern access to the transmission medium, and an interface between the station and transmission medium. A software element is required to implement the medium access protocol for intercommunication between stations and for preparation of frames for transmission, and vice versa. These hardware and software elements, all of which are implemented at the lowest two layers of the

OSI Reference Model, are usually combined into a **Network Interface Card (NIC)**, also known as a **Network Adapter**. An NIC interfaces with a standard computer such as a PC and at the physical layer connects to the particular transmission media to be used. Another software element is therefore also required to regulate the interface between the computer system and NIC.

9.1**Medium Access Control**

The media used in LANs generally convey frames from only one station at a time, although the media themselves are generally shared by a number of stations. In order to overcome the difficulties which may arise through sharing, a **Medium Access Control (MAC)** mechanism or protocol is necessary. A MAC protocol merely regulates how stations may access the medium in an orderly fashion for correct operation and also attempts to ensure that each station obtains a fair share of its use.

LAN networks usually have only a single medium over which all messages, represented over a series of frames, are transmitted. If the medium is not being used two, or more, stations may simultaneously attempt an access, leading to a collision. An MAC technique is therefore required to regulate access by stations to the medium and handle the effect of two, or more, stations simultaneously attempting to access the medium. There is also the danger that once a pair of stations have established communication, all other stations may be excluded, perhaps indefinitely, or at least for a considerable period of time.

A LAN does not usually have any separate network control function for operation. Nor is a separate control function required to detect abnormal network conditions, or to control recovery therefrom. Rather, each station is generally equally responsible in a LAN, in which case control is said to be **fully distributed**.

Three general MAC techniques exist for use within fully distributed networks:

1. **Contention:** Here there is no regulating mechanism directly to govern stations attempting to access a medium. Rather, two or more stations may contend for the medium and any multiple simultaneous accesses are resolved as they arise.
2. **Token passing:** A single **token** exists within the network and is passed between stations in turn. Only a station holding the token may use the medium for transmission. This eliminates multiple simultaneous accesses of the medium with the attendant risk of collision.
3. **Slotted and register insertion rings:** Similar in principle to token passing, but a unique time interval is granted to a station for transmission.

Carrier Sense Multiple Access

Many modern LANs evolved from a LAN known as **Aloha** which was one of the first primitive LANs to be developed. Aloha was packet based and used radio as its transmission medium. It was used by the University of Hawaii in the early 1970s to inter-

JTAM Job transfer, access and management.

Link management A function of the data link layer of the OSI Reference Model which is concerned with setting up and disconnection of a link.

Local area network (LAN) A data communications network used to interconnect a community of digital devices distributed over a localized area of up to, say, 10 km². The devices may be office workstations, mini- and microcomputers, intelligent instrumentation equipment, and so on.

Logical Link Control (LLC) A protocol forming part of the data link layer in LANs. It is concerned with the reliable transfer of data across the data link between two communicating systems.

Management information base (MIB) The name of the database used to hold the management information relating to a network or internetwork.

Manchester encoding A 1B2B code which converts each single binary 1 and binary 0 into two respective equal, and opposite, binary signal elements.

Mark A term traditionally used in telegraph systems to indicate a logic 1 state of a bit.

Maximum burst size (MBS) A traffic parameter, agreed between an ATM network and a customer, that is the agreed maximum size of a burst of cells with variable bit rate that can be accepted by the network.

Media gateway A device that converts media provided in one type of network to the format required for another type of network.

Media Gateway Control Protocol A protocol for controlling media gateways.

Medium access control (MAC) A method of determining which device has access to a shared transmission medium in a local area network.

Metropolitan area network (MAN) A network that links a set of LANs that are physically distributed around a town or city.

Modem The device that converts a binary (digital) data stream into an analogue (continuously varying) form, prior to transmission of the data across an analogue network (MODulator), and reconverts the received signal back into its binary form (DEModulator). Since each access port to the network normally requires a full-duplex (two-way simultaneous) capability, the device must perform both the MODulation and the DEModulation functions; hence the single name MODEM is used. As an example, a modem is normally required to transmit data across a telephone network.

Moving Picture Experts Group (MPEG) An ISO committee that generates standards for digital video compression. It also gives its name to their standards.

Multidrop A type of network configuration that supports more than two stations on the same transmission medium.

Multiplexer A device to enable a number of lower bit rate devices, normally situated in the same location, to share a single higher bit rate transmission line. The data-carrying capacity of the latter must be in excess of the combined bit rates of the low bit rate devices.

Multi Protocol Label Switching A standard that integrates layer 3 routing using IP addresses with a layer 2 switching technique such as ATM or Frame Relay.

Network interface card A physical interface in an end system such as a computer that connects to a transmission medium.

Network layer This corresponds to layer 3 of the ISO reference model for open systems interconnection. It is concerned with the establishment and clearing of logical or physical connections across the network being used.

Appendix D

Comer

CD-ROM by
Droms

Computer Networks and Internets

**Second
Edition**



**TECH
300
C664
1999**

Computer Networks And Internets

Second Editon

Douglas E. Comer

Department of Computer Sciences
Purdue University
West Lafayette, IN 47907



PRENTICE HALL
Upper Saddle River, New Jersey 07458

Library of Congress Cataloging-in-Publications Data

Comer, Douglas.

Computer networks and internets / Douglas E. Comer, -- 2nd ed.

p. cm.

Includes index.

ISBN 0-13-083617-6

1. Computer networks. 2. Internetworking (Telecommunication)

I. Title.

TK5105.5.C5897 1999

004.6--dc21

98-47110

CIP

Publisher: **ALAN APT**

Development Editor: **SONDRA CHAVEZ**

Editor-in-chief: **MARCIA HORTON**

Production editor: **MIMI JETT—Interactive Composition Corporation**

Managing editor: **EILEEN CLARK**

Director of production and manufacturing: **DAVID W. RICCARDI**

Cover director: **HEATHER SCOTT**

Manufacturing buyer: **DONNA M. SULLIVAN**

Manufacturing manager: **TRUDY PISCIOOTTI**

Editorial assistant: **TONI HOLM**



© 1999 by Prentice-Hall, Inc.
Simon & Schuster / A Viacom Company
Upper Saddle River, New Jersey 07458

All rights reserved. No part of this book may be reproduced, in any form or by any means, without permission in writing from the publisher.

The author and publisher of this book have used their best efforts in preparing this book. These efforts include the development, research, and testing of the theories and programs to determine their effectiveness. The author and publisher make no warranty of any kind, expressed or implied, with regard to these programs or the documentation contained in this book. The author and publisher shall not be liable in any event for incidental or consequential damages in connection with or arising out of the furnishing, performance, or use of these programs.

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

ISBN 0-13-083617-6

Prentice-Hall International (UK) Limited, London
Prentice-Hall of Australia Pty. Limited, Sydney
Prentice-Hall Canada Inc., Toronto
Prentice-Hall Hispanoamericana, S.A., Mexico
Prentice-Hall of India Private Limited, New Delhi
Prentice-Hall of Japan, Inc., Tokyo
Simon & Schuster Asia Pte. Ltd., Singapore
Editora Prentice-Hall do Brasil, Ltda., Rio de Janeiro

UNIX is a registered trademark of UNIX Systems

Laboratories, Incorporated

Microsoft Windows is a trademark of Microsoft Corporation

Microsoft is a trademark of Microsoft Corporation

Windows 95 is a trademark of Microsoft Corporation

Windows 98 is a trademark of Microsoft Corporation

Windows NT is a trademark of Microsoft Corporation

Solaris is a trademark of Sun Microsystems, Incorporated

Sniffer is a trademark of Network General Corporation

Java is a trademark of Sun Microsystems, Incorporated

JavaScript is a trademark of Sun Microsystems, Incorporated

AdaMagic is a trademark of Intermetrics, Incorporated

Sparc is a trademark of Sun Microsystems, Incorporated

Alpha is a trademark of Digital Equipment Corporation

Pentium is a trademark of Intel Corporation

X Window System is a trademark of X Consortium, Inc.

7.3 Shared Communication Channels

The history of computer networking changed dramatically during the late 1960s and early 1970s when researchers developed a form of computer communication known as *Local Area Networks (LANs)*. Devised as alternatives to expensive, dedicated point-to-point connections, the designs differ fundamentally from long-distance networks because they rely on sharing the network. Each LAN consists of a single shared medium, usually a cable, to which many computers attach. The computers take turns using the medium to send packets.

Several LAN designs emerged from the research. The designs differ in details such as the voltages and modulation techniques used, and the approach to sharing (i.e., the mechanisms used to coordinate access and transmit packets).

Because it eliminates duplication, sharing has an important economic impact on networking: it reduces cost. Consequently, Local Area Network technologies that allow a set of computers to share a medium have become popular. In fact,

Networks that allow multiple computers to share a communication medium are used for local communication. Point-to-point connections are used for long-distance networks and a few other special cases.

If sharing reduces cost, why are shared networks used only for local communication? Both technical and economic reasons contribute to the answer. We said that the computers attached to a shared network must coordinate use of the network. Because coordination requires communication and the time required to communicate depends on distance, a large geographic separation between computers introduces longer delays. Thus, shared networks with long delays are inefficient because they spend more time coordinating use of the shared medium and less time sending data. In addition, engineers have learned that providing a high bandwidth communication channel over long distances is significantly more expensive than providing the same bandwidth communication over a short distance.

7.4 Significance Of LANs And Locality Of Reference

The significance of LANs can be stated simply:

Local Area Network technologies have become the most popular form of computer networks. LANs now connect more computers than any other type of network.

One of the reasons so many LANs have been installed is economic: LAN technologies are both inexpensive and widely available. However, the main reason the demand for LANs is high can be attributed to a fundamental principle of computer networking known as *locality of reference*. The locality of reference principle states that commun-

cation among a set of computers is not random, but instead follows two patterns. First, if a pair of computers communicates once, the pair is likely to communicate again in the near future and then periodically. The pattern is called *temporal locality of reference* to imply a relationship over time. Second, a computer tends to communicate most often with other computers that are nearby. The second pattern is called *physical locality of reference*[†] to emphasize the geographic relationship. We can summarize:

The locality of reference principle: *computer communication follows two distinct patterns. First, a computer is more likely to communicate with computers that are physically nearby than with computers that are far away. Second, a computer is more likely to communicate with the same set of computers repeatedly.*

The locality of reference principle is easy to understand because it applies to human communication. For example, people communicate most often with others who are physically nearby (e.g., working together). Furthermore, if an individual communicates with someone (e.g., a friend or family member), the individual is likely to communicate with the same person again.

7.5 LAN Topologies

Because many LAN technologies have been invented, it is important to know how specific technologies are similar and how they differ. To help understand similarities, each network is classified into a category according to its *topology* or general shape. This section describes the three topologies used most often with LANs; later sections add more detail and show specific examples.

7.5.1 Star Topology

A network uses a *star topology* if all computers attach to a central point. Figure 7.3 illustrates the concept.

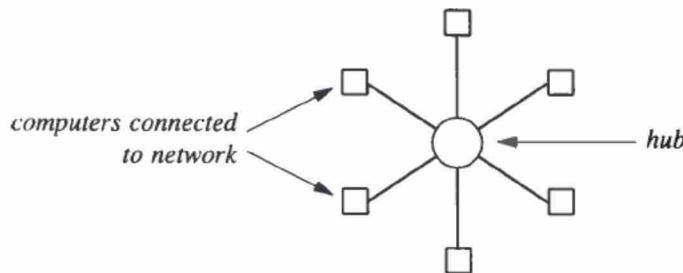


Figure 7.3 Illustration of the star topology in which each computer attaches to a central point called a *hub*.

[†]Physical locality of reference is sometimes referred to as *spatial locality of reference*.

Because a star-shaped network resembles the spokes of a wheel, the center of a star network is often called a *hub*. A typical hub consists of an electronic device that accepts data from a sending computer and delivers it to the appropriate destination.

Figure 7.3 illustrates an idealized star network. In practice, star networks seldom have a symmetric shape in which the hub is located an equal distance from all computers. Instead, a hub often resides in a location separate from the computers attached to it. For example, Chapter 9 will illustrate that computers can reside in individual offices, while the hub resides in a location accessible to an organization's networking staff.

7.5.2 Ring Topology

A network that uses a *ring topology* arranges for computers to be connected in a closed loop – a cable connects the first computer to a second computer, another cable connects the second computer to a third, and so on, until a cable connects the final computer back to the first. The name *ring* arises because one can imagine the computers and the cables connecting them arranged in a circle as Figure 7.4 illustrates.

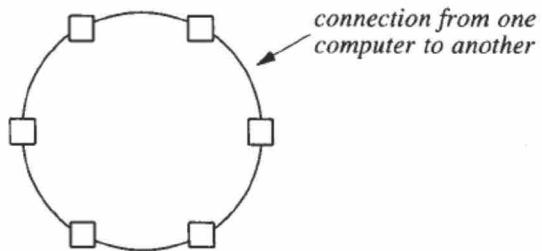


Figure 7.4 Illustration of a ring topology in which computers are connected in a closed loop. Each computer connects directly to two others.

It is important to understand that the *ring*, like the star topology, refers to logical connections among computers, not physical orientation – the computers and connections in a ring network need not be arranged in a circle. Instead, the cable between a pair of computers in a ring network may follow a hallway or rise vertically from one floor of a building to another. Furthermore, if one computer is far from others in the ring, the two cables that connect the distant computer may follow the same physical path.

7.5.3 Bus Topology

A network that uses a *bus topology* usually consists of a single, long cable to which computers attach†. Any computer attached to a bus can send a signal down the cable, and all computers receive the signal. Figure 7.5 illustrates the topology. Because all computers attached to the cable can sense an electrical signal, any computer can send

†In practice, the ends of a bus network must be terminated to prevent electrical signals from reflecting back along the bus.

data to any other computer. Of course, the computers attached to a bus network must coordinate to ensure that only one computer sends a signal at any time or chaos results.

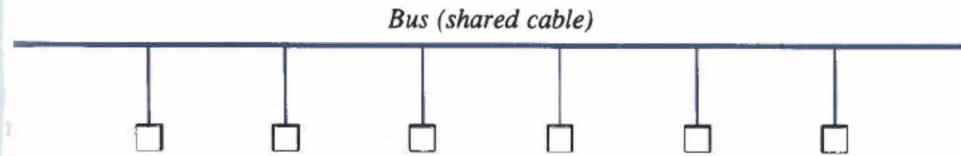


Figure 7.5 Illustration of a bus topology in which all computers attach to a single cable.

7.5.4 The Reason For Multiple Topologies

Each topology has advantages and disadvantages. A ring topology makes it easy for computers to coordinate access and to detect whether the network is operating correctly. However, an entire ring network is disabled if one of the cables is cut. A star topology helps protect the network from damage to a single cable because each cable connects only one machine. A bus requires fewer wires than a star, but has the same disadvantage as a ring: a network is disabled if someone accidentally cuts the main cable. In addition to later sections in this chapter, other chapters provide detailed examples of network technologies that illustrate some of the differences.

We can summarize the major points about network topologies.

Networks are classified into broad categories according to their general shape. The primary topologies used with LANs are star, ring, and bus; each topology has advantages and disadvantages.

7.6 Example Bus Network: Ethernet

7.6.1 History Of The Ethernet

Ethernet is a well-known and widely used network technology that employs bus topology. Ethernet was invented at Xerox Corporation's Palo Alto Research Center in the early 1970s. Digital Equipment Corporation, Intel Corporation, and Xerox later cooperated to devise a production standard, which is informally called *DIX Ethernet* for the initials of the three companies. IEEE now controls Ethernet standards†. In its original version, an Ethernet LAN consisted of a single coaxial cable, called the *ether*, to which multiple computers connect. Engineers use the term *segment* to refer to the Ethernet coaxial cable. A given Ethernet segment is limited to 500 meters in length, and the standard requires a minimum separation of 3 meters between each pair of connections.

†Several variations of Ethernet currently exist; this section describes the original technology and leaves the discussion of alternatives until Chapter 9.

The original Ethernet hardware operated at a bandwidth of 10 Megabits per second (Mbps); a later version known as *Fast Ethernet* operates at 100 Mbps, and the most recent version, which is known as *Gigabit Ethernet* operates at 1000 Mbps or 1 Gigabit per second (Gbps).

7.6.2 Sharing On An Ethernet

The Ethernet standard specifies all details, including the format of frames that computers send across the ether†, the voltage to be used, and the method used to modulate a signal.

Because it uses a bus topology, Ethernet requires multiple computers to share access to a single medium. A sender transmits a signal, which propagates from the sender toward both ends of the cable. Figure 7.6 illustrates how data flows across an Ethernet.

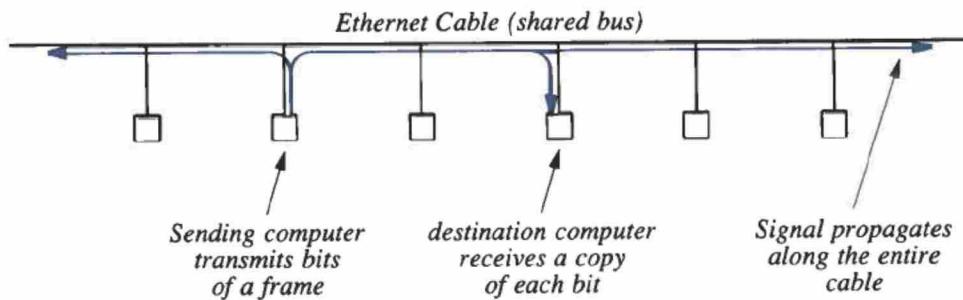


Figure 7.6 Conceptual flow of bits across an Ethernet. While transmitting a frame, a computer has exclusive use of the cable.

As the figure shows, a signal propagates from the sending computer to both ends of the shared cable. It is important to understand that sharing in local area networks technologies, does not mean that multiple frames are being sent at the same time. Instead, the sending computer has exclusive use of the entire cable during the transmission of a given frame – other computers must wait. After one computer finishes transmitting one frame, the shared cable becomes available for another computer to use. To summarize:

Ethernet is a bus network in which multiple computers share a single transmission medium. While one computer transmits a frame to another, all other computers must wait.

†Chapter 8 discusses Ethernet frames in more detail and shows an example.

7.7 Carrier Sense On Multi-Access Networks (CSMA)

The most interesting aspect of Ethernet is the mechanism used to coordinate transmission. An Ethernet network does not have a centralized controller that tells each computer how to take turns using the shared cable. Instead, all computers attached to an Ethernet participate in a distributed coordination scheme called *Carrier Sense Multiple Access (CSMA)*. The scheme uses electrical activity on the cable to determine status. When no computer is sending a frame, the ether does not contain electrical signals. During frame transmission, however, a sender transmits electrical signals used to encode bits. Although the signals differ slightly from the carrier waves described in Chapter 5, they are informally called a *carrier*. Thus, to determine whether the cable is currently being used, a computer can check for a carrier. If no carrier is present, the computer can transmit a frame. If a carrier is present, the computer must wait for the sender to finish before proceeding. Technically, checking for a carrier wave is called *carrier sense*, and the idea of using the presence of a signal to determine when to transmit is called *Carrier Sense Multiple Access (CSMA)*.

7.8 Collision Detection And Backoff With CSMA/CD

Because CSMA allows each computer to determine whether a shared cable is already in use by another computer, it prevents a computer from interrupting an ongoing transmission. However, CSMA cannot prevent all possible conflicts. To understand why, imagine what happens if two computers at opposite ends of an idle cable both have a frame ready to send at the same time. When they check for a carrier, both stations find the cable idle, and both start to send frames simultaneously. The signals travel at approximately 70% of the speed of light, and when the signals transmitted by two computers reach the same point on the cable, they interfere with each other.

The interference between two signals is called a *collision*. Although a collision does not harm the hardware, it produces a garbled transmission that prevents either of the two frames from being received correctly. To ensure that no other computer transmits simultaneously, the Ethernet standard requires a sending station to monitor signals on the cable. If the signal on the cable differs from the signal that the station is sending, it means that a collision has occurred[†]. Whenever a collision is detected, a sending station immediately stops transmitting. Technically, monitoring a cable during transmission is known as *Collision Detect (CD)*, and the Ethernet mechanism is known as *Carrier Sense Multiple Access with Collision Detect (CSMA/CD)*.

CSMA/CD does more than merely detect collisions – it also recovers from them. After a collision occurs, a computer must wait for the cable to become idle again before transmitting a frame. However, if the computers begin to transmit as soon as the ether becomes idle, another collision will occur. To avoid multiple collisions, Ethernet requires each computer to delay after a collision before attempting to retransmit. The standard specifies a maximum delay, d , and forces each computer to choose a random delay less than d . In most cases, when a computer chooses a delay at random, it will

[†]To guarantee that a collision has time to reach all stations before they stop transmitting, the Ethernet standard specifies both a maximum cable length and a minimum frame size.

select a value that differs from any of the values chosen by the other computers – the computer that chooses the smallest delay will proceed to send a frame and the network will return to normal operation.

If two or more computers happen to choose nearly the same amount of delay after a collision, they will both begin to transmit at nearly the same time, producing a second collision. To avoid a sequence of collisions, Ethernet requires each computer to double the range from which a delay is chosen after each collision. Thus, a computer chooses a random delay from 0 to d after one collision, a random delay between 0 and $2d$ after a second collision, between 0 and $4d$ after a third, and so on. After a few collisions, the range from which a random value is chosen becomes large, and the probability is high that some computer will choose a short delay and transmit without a collision.

Technically, doubling the range of the random delay after each collision is known as *binary exponential backoff*. In essence, exponential backoff means that an Ethernet can recover quickly after a collision because each computer agrees to wait longer times between attempts when the cable becomes busy. In the unlikely event that two or more computers choose delays that are approximately equal, exponential backoff guarantees that contention for the cable will be reduced after a few collisions. We can summarize:

Computers attached to an Ethernet use CSMA/CD in which a computer waits for the ether to be idle before transmitting a frame. If two computers transmit simultaneously, a collision occurs; the computers use exponential backoff to choose which computer will proceed. Each computer delays a random time before trying to transmit again, and then doubles the delay for each successive collision.

7.9 Wireless LANs And CSMA/CA

A set of *wireless LAN* technologies are available that use a modified form of CSMA/CD. The products, which are manufactured by several companies are available under a variety of trade names. For example, NCR Corporation sells *WaveLAN*, Solectek sells *AirLAN*, and Proxim Corporation sells *RangeLAN*.

Instead of transmitting signals across a cable, wireless LAN hardware uses antennas to broadcast RF signals through the air, which other computers receive. The devices use 900 MHz frequencies to permit data to be sent at 2 Mbps. Like other LAN technologies, the wireless LANs use sharing. That is, all the computers participating in a given wireless LAN are configured to the same radio frequency. Thus, they must take turns sending packets.

One difference between the way wired and wireless LANs manage sharing arises because of the way wireless transmissions propagate. Although the electromagnetic energy radiates in all directions, wireless LAN transmitters use low power, meaning that a transmission only has enough power to travel a short distance. Furthermore, metallic obstructions can block the signal. Thus, wireless units located far apart or behind obstructions will not receive each other's transmissions.

easily carry data in two directions simultaneously, each connection uses a pair of fibers as Figure 7.11 illustrates.

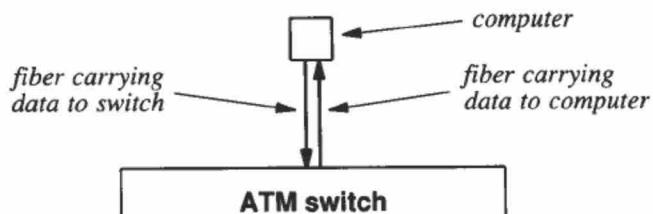


Figure 7.11 Details of a connection between an ATM switch and a computer.

Each connection consists of a pair of optical fibers. One fiber carries data to the switch, and the other carries data to the computer.

Like the optical fibers used with FDDI, the pair of fibers used to connect a computer to an ATM switch are fastened together. Usually, the jacket on one fiber contains a colored stripe or is labeled; whoever installs a connection uses the label to ensure that the output of the switch connects to the input of the computer and vice versa.

To summarize:

An ATM network is formed from a switch to which multiple computers attach. The connection between a computer and an ATM switch consists of a pair of fibers, one carrying data in each direction.

7.14 Summary

This chapter discusses an alternative to direct point-to-point communication called a Local Area Network (LAN). Designed for use over a small distance (e.g., in a building), a LAN does not need a separate wire between each pair of computers. Instead, a LAN consists of a single, shared medium to which many computers attach. The computers take turns using the medium to send data.

Although LAN technologies require computers to divide data into small packets called frames, only one packet can be transmitted on a LAN at any time. That is, while transmitting, a computer has exclusive use of the LAN. To make access fair, each computer is permitted to hold the shared medium for the transmission of one frame before allowing another computer to proceed. Thus, after it gains control, a computer sends a frame and then relinquishes control to another computer.

Each computer network can be classified into one of a few basic categories, depending on its topology. A bus topology consists of a single, shared cable to which many computers attach. When it uses a bus, a computer transmits a signal that all other computers attached to the bus receive. A ring topology consists of computers connected in a closed loop. The first computer connects to the second, the second connects to the third, and so on, until the last computer connects back to the first. Finally, a star topology resembles a wheel with the network itself corresponding to a central hub, and the links to individual computers corresponding to spokes. Each topology has advantages and disadvantages; no topology is best for all purposes.

LAN technologies exist that use each topology. An Ethernet LAN uses a bus topology, as does LocalTalk. To access an Ethernet, stations obey Carrier Sense Multiple Access with Collision Detect (CSMA/CD). That is, a station waits for the ether to be idle, and then attempts to send. If two stations transmit at the same time, a collision results, causing them to wait a random time before trying again. Successive collisions cause exponential backoff in which each station doubles its delay.

Wireless LANs such as WaveLAN, RangeLAN, or AirLAN use Carrier Sense Multiple Access With Collision Avoidance (CSMA/CA). Before transmitting a data frame, a sender transmits a small control message to which the receiver responds. The exchange of control messages notifies all stations within range of the receiver that a data transmission is about to occur. Other stations then remain silent while the transmission takes place (i.e., avoid a collision), even if they do not receive a copy of the signal.

Stations attached to a token passing ring network also share the medium. While one station transmits a frame, all other stations pass the bits around the ring, which allows the sender to verify that the bits were transmitted correctly. To coordinate use of the ring and guarantee fairness, stations on a token ring send a special message called a token. A station waits for the token to arrive, uses the complete ring to transmit one frame, and then sends the token to the next station. IBM Token Ring and FDDI networks both use token passing. FDDI differs from conventional token passing technologies because it can be configured with an extra ring that is used to recover from catastrophic failures. The extra ring is called counter-rotating because data flows the opposite direction than on the main ring. An FDDI network with a counter-rotating ring is said to be self-healing because it can detect a failure and loop back along the reverse ring to close the path.

LANs that use ATM technology have a star topology. An ATM switch forms the hub of the star to which each computer connects. Because ATM is designed to operate at high speed, the connection between a computer and an ATM switch uses a pair of optical fibers, with one fiber carrying data in each direction.

Appendix E

WEBSTER'S NEW WORLD™



TELECOM DICTIONARY

*A comprehensive reference for
telecommunications terminology*

*More than 4,600 telecommunications terms and
acronyms defined and explained*

*Clear, up-to-date coverage of recent data communications
terminology, including VoIP, SANs, 802.11n, and Bluetooth*

RAY HORAK

WE DEFINE YOUR WORLD®

WEBSTER'S NEW WORLDTM

Telecom Dictionary

Ray Horak



Wiley Publishing, Inc.

Webster's New World® Telecom Dictionary

Published by
Wiley Publishing, Inc.
10475 Crosspoint Boulevard
Indianapolis, IN 46256
www.wiley.com

Copyright © 2008 by Wiley Publishing, Inc., Indianapolis, Indiana

Published simultaneously in Canada

ISBN: 978-0-471-77457-0

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

Library of Congress Cataloging-in-Publication Data

Horak, Ray.

Webster's New World telecom dictionary / Ray Horak.

p. cm.

ISBN 978-0-471-77457-0 (pbk.)

1. Telecommunication—Dictionaries. I. Title.

TK5102.H65 2007

621.38203—dc22

2007024232

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600. Requests to the Publisher for permission should be addressed to the Legal Department, Wiley Publishing, Inc., 10475 Crosspoint Blvd., Indianapolis, IN 46256, (317) 572-3447, fax (317) 572-4355, or online at <http://www.wiley.com/go/permissions>.

Limit of Liability/Disclaimer of Warranty: The publisher and the author make no representations or warranties with respect to the accuracy or completeness of the contents of this work and specifically disclaim all warranties, including without limitation warranties of fitness for a particular purpose. No warranty may be created or extended by sales or promotional materials. The advice and strategies contained herein may not be suitable for every situation. This work is sold with the understanding that the publisher is not engaged in rendering legal, accounting, or other professional services. If professional assistance is required, the services of a competent professional person should be sought. Neither the publisher nor the author shall be liable for damages arising herefrom. The fact that an organization or Website is referred to in this work as a citation and/or a potential source of further information does not mean that the author or the publisher endorses the information the organization or Website may provide or recommendations it may make. Further, readers should be aware that Internet Websites listed in this work may have changed or disappeared between when this work was written and when it is read.

For general information on our other products and services please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Trademarks: Wiley, the Wiley logo, Webster's New World, the Webster's New World logo, We Define Your World, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates, in the United States and other countries, and may not be used without written permission. All other trademarks are the property of their respective owners. Wiley Publishing, Inc., is not associated with any product or vendor mentioned in this book.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

dialing the destination telephone number. Synonymous with *caller ID*. See also *ACD*, *ANI*, *call center*, *CLASS*, *LEC*, *PSTN*, and *screen pop*.

2. A voice telephone system feature that supports the CLID network service and offers a similar capability for station-to-station PBX calls. See also *CLASS*.

client In a client/server architecture, a complete, standalone computer that optimizes the user interface, relying on servers to handle the more mundane tasks associated with application and file storage, network administration, security, and other critical functions. See also *architecture*, *client/server*, and *server*.

client mesh See *pure mesh*.

client/server A network architecture that distributes intelligence and responsibilities at several levels, with some machines designated as servers to serve the needs of client machines. A server can be a mainframe, minicomputer, or personal computer that operates in a time-sharing mode to provide for the needs of many clients. Client machines are complete, standalone computers that optimize the user interface, relying on servers to handle the more mundane tasks associated with application and file storage, network administration, security, and other critical functions. See also *peer-to-peer*.

Clipper Chip An integrated circuit that uses the Skipjack voice encryption algorithm developed by the United States National Security Agency (NSA) for the National Institute of Science and Technology (NIST). Skipjack is a block coding algorithm that encrypts 64-bit data blocks with an 80-bit key. Data encrypted by the Skipjack algorithm can be provided not only to the intended recipient through the use of a key, but also by the U.S. government through the use of a back door into a Law Enforcement Access Field (LEAF). The Clipper Chip is manufactured by the U.S. government, which has tried unsuccessfully to make it, and similar technologies, mandatory for voice encryption in the United States. Privacy advocates feared that government authorities would abuse the back door. Law enforcement authorities fear that the widespread use of other voice encryption technologies will make it impossible to place legal wiretaps. See also *algorithm*, *back door*, *encryption*, *integrated circuit*, and *wiretap*.

CLNP (ConnectionLess Network Protocol) A Network Layer datagram protocol from the International Organization for Standardization (ISO) for use over OSI (Open Systems Integration) networks and specified in ISO 8473. CLNP is very similar to Internet Protocol (IP). The datagram size is the same as IP, and there are similar mechanisms for fragmentation, error control, and lifetime control. CLNP, however, has an address space of 20 octets compared the IPv4 address space of only 4 octets. OSI networks have not been well accepted, however, and the OSI protocol stack has been relegated to the status of OSI Reference Model. See also *datagram*, *error control*, *fragmentation*, *IP*, *ISO*, *lifetime control*, *Network Layer*, *OSI*, *OSI Reference Model*, *protocol*, and *protocol stack*.

clocking pulse Periodic signals generated by a timing source for purposes of synchronizing the flow of data within a computer or between computers across a circuit. See also *synchronous transmission*.

closed circuit television (CCTV) See *CCTV*.

closed-loop algorithm In frame relay, a congestion control mechanism that prevents the frame relay network device (FRND) from accepting incoming frames unless there is an extremely high probability of the network's being able to deliver them without discard. A closed-loop algorithm fairly allocates backbone bandwidth among all the permanent virtual circuits (PVCs) configured on a particular trunk, and in proportion to the Committed Information Rate (CIR) of each PVC. See also *backbone*, *bandwidth*, *CIR*, *congestion*, *frame relay*, *FRND*, *PVC*, and *trunk*.

closed user group (CUG) See *CUG*.

cloud A wide area network (WAN) commonly is depicted as a cloud, which serves to obscure its complex inner workings from view. Data just pops in on one side of the cloud and pops out on the other side, so to speak.

LAN (Local Area Network) A LAN is a packet network designed to interconnect host computers, peripherals, storage devices, and other computing resources within a local area, i.e., limited distance. LANs conform to the client/server architecture, a distributed computing architecture that runs applications on client microcomputers against one or more centralized servers, which are high-performance multiport computers with substantial processing power and large amounts of memory. A LAN might serve an office, a floor of a building, and entire building, or a campus area, but generally does not cross a public right-of-way such as a street. The distance limitation generally is in the range of a few kilometers, at most, although that is sensitive to the transmission media employed, which include coaxial cable, twisted pair, optical fiber, infrared (IR) light, and radio frequency (RF) systems. Raw bandwidth ranges up to 10 Gbps, although actual throughput generally is much less. LANs generally are private networks, although public wireless hotspots offering wireless Internet access currently are popular. Most LAN standards are set by the 802 Working Group of the Institute of Electrical and Electronic Engineers (IEEE), with examples being 802.3 (Ethernet) and 802.11a/b/g (Wi-Fi). A personal area network (PAN) such as Bluetooth, is much more limited in geographic scope than a LAN. LANs and LAN segments can be interconnected over a metropolitan area network (MAN) or wide area network (WAN). LANs operate at Layer 1, the Physical Layer, and Layer 2, the Data Link Layer, of the OSI Reference Model. See also *802.3, 802.5, 802.11, architecture, bandwidth, Bluetooth, client/server, coaxial cable, Data Link Layer, Ethernet, hotspot, IEEE, IR, MAN, optical fiber, OSI Reference Model, PAN, Physical Layer, RF, throughput, Token Ring, twisted pair, WAN, and Wi-Fi*.

landline 1. A traditional telephone connected to the PSTN by a traditional wire (or fiber) local loop that terminates in a fixed location, rather than a cellular mobile telephone connected to a cellular network via radio technology. A cordless telephone is considered part of a landline as the local loop terminates in a fixed base station on the subscriber premises, even though the connection to the base station is wireless. A wireless local loop (WLL) is considered a landline, as it is terrestrial and connects two fixed points. See also *local loop and WLL*. 2. A telecommunications system that uses traditional terrestrial cabled, or conducted, transmission media such as copper or fiber optics, and wireless systems such as microwave, rather than mobile wireless radio technologies such as cellular or, especially, non-terrestrial satellite.

LANE (LAN Emulation) A specification (January 1995) from the ATM Forum (since merged into the MFA Forum) for an ATM service in support of native Ethernet (802.3) and Token Ring (802.5) local area network (LAN) communications over an ATM network. Software in the end systems (e.g., ATM-based hosts or routers, known as *proxies*), of the ATM network emulates a native LAN environment. LANE acts as Layer 2 bridge in support of connectionless LAN traffic, with the connection-oriented ATM service being transparent to the user application. In LANE, a LAN emulation client (LEC) connects to the ATM network over a LANE user-to-network interface (LUNI). The network-based LAN emulation server (LES) registers the LAN medium access control (MAC) addresses and translates them into ATM addresses using the address resolution protocol (ARP). Each LEC is assigned to an emulated LAN (ELAN) by an optional network-based LAN emulation configuration server (LECS). Each LEC also is associated with a broadcast and unknown server (BUS) that handles broadcast and multicast traffic, as well as initial unicast frames before address resolution. LANE traffic generally is Class C variable bit rate (VBR) traffic in message mode, and is supported over ATM Adaptation Layer Type 5 (AAL5). See also *802.3, 802.5, AAL5, ARP, ATM, ATM Forum, broadcast, BUS, Class C ATM traffic, connectionless, connection-oriented, ELAN, emulation, Ethernet, host, Layer 2, LEC, LECS, LES, LUNI, MAC, message mode service, MFA Forum, multicast, proxy, router, Token Ring, unicast, and VBR*.

LAN emulation (LANE) See *LANE*.

LAN emulation client (LEC) See *LEC*.

LAN emulation configuration server (LECS) See *LECS*.

LAN emulation server (LES) See *LES*.

LANE user-to-network interface (LUNI) See *LUNI*.

Appendix F



VOIP Industry Dictionary is a free tool to get technical definitions and lookup industry terms from more than 11,000 terms, acronyms, and definitions.

Dictionaries:

[» Telecom](#)

[» CATV](#)

[» eMarketing](#)

[» Optical](#)

[» IPTV](#)

[» VOIP](#)

[» Conatct Us](#)



» Data Network

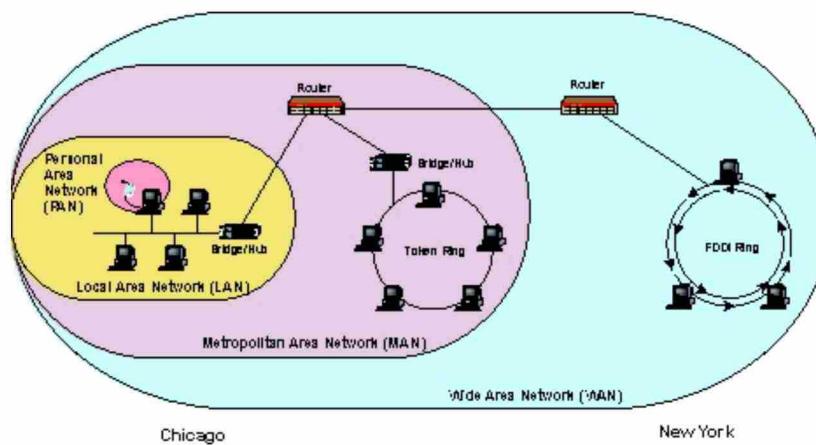
A data network is a system that **transfers data** between network access points (nodes) **through data switching**, system control and interconnection transmission lines. Data networks are primarily designed to transfer data from one point to one or more points (multipoint). Data networks may be composed of a variety of communication systems including circuit switches, leased lines and packet switching networks. There are predominately two types of data networks, broadcast and point-to-point.

Ruckus Networks

Patchsee is a Ethernet patching solution that is unique and provides efficiency.

Wisyne

Open



Data Network Type Diagram

[Lookup](#)

This industry dictionary is powered by [Althos Telecom Dictionary](#)

» Add this Online Telecom Dictionary tool to your web site

In order to add this online TeleCom industry dictionary tool to your web site and give your visitors the way to lookup industry terms from more than 11,000 terms, acronyms, and definitions, just copy the HTML code below and put it into your web page where you want to provide access to your dictionary.

```

<table align="center" cellspacing="1" cellpadding="5" border="0"
bgcolor="#C0C0C0">
<tr bgcolor="#FFFFFF"><td><font style="font-size: 11px; font-
family: Verdana;"><b>Find TeleCom industry <br>terms,
acronyms, and definitions instantly:</b></font></td></tr>
<FORM METHOD="POST">

```

Other Dictionaries: [Wireless Dictionary](#) | [Optical Dictionary](#) | [Billing Dictionary](#) | [Telecom Dictionary](#) | [CATV Dictionary](#) | [eMarketing Dictionary](#) | [IPTV Dictionary](#) | [VOIP Dictionary](#)

Related Websites: [Althos](#) - [Althos Books](#) - [Telecom Books](#) - [Telecom Dictionary](#)

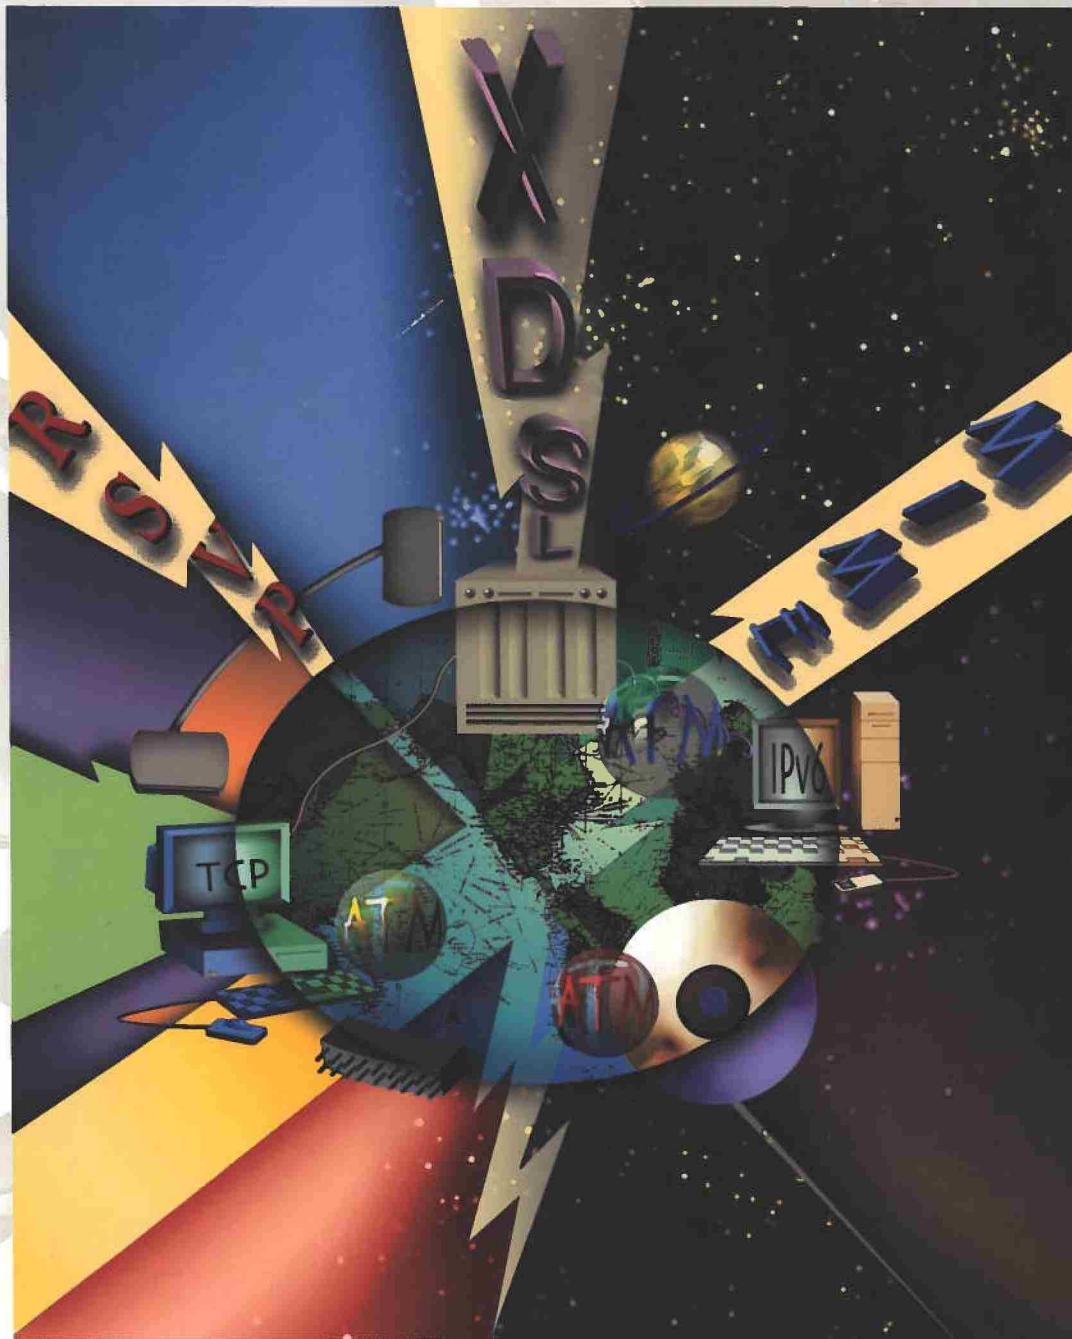
Copyright  2009 All rights reserved

Disclaimer: Telecom Dictionary Tool is the trademarks of Althos. Althos is not affiliated with any other website and provides the information in this dictionary from content created by our expert writers. Althos provides our services on "as is" and "as available" basis and we do not provide any guarantees regarding this service's accuracy and/or availability.

Appendix G

DATA & COMPUTER COMMUNICATIONS

SIXTH EDITION



WILLIAM STALLINGS

DATA AND COMPUTER COMMUNICATIONS

SIXTH EDITION

William Stallings

Business Library
University of Wisconsin-Madison
2200 Grainger Hall
975 University Avenue
Madison WI 53706-1323



Prentice Hall
Upper Saddle River, New Jersey 07458

Library of Congress Cataloging-in-Publication Data

Stallings, William.

Data and computer communications / William Stallings.—6th ed.

p. cm.

Includes bibliographical references and index.

ISBN 0-13-084370-9

1. Data transmission systems. 2. Computer networks. I. Title.

TK5105.S73 1999

004.6—dc20

99-24419

CIP

Publisher: Alan Apt

Editor-in-chief: Marcia Horton

Project manager: Ana Arias Terry

Editorial/production supervision: Rose Kernan

Executive managing editor: Vince O'Brien

Copy editing: Patricia Daly

Art director and cover designer: Heather Scott

Assistant Vice-president of production: David W. Riccardi

Manufacturing buyer: Pat Brown

Editorial assistant: Toni Holm

© 2000, 1996 by Prentice-Hall, Inc.
Upper Saddle River, New Jersey 07458

The author and publisher of this book have used their best efforts in preparing this book. These efforts include the development, research, and testing of the theories and programs to determine their effectiveness. The author and publisher make no warranty of any kind, expressed or implied, with regard to these programs or the documentation contained in this book. The author and publisher shall not be liable in any event for incidental or consequential damages in connection with, or arising out of, the furnishing, performance, or use of these programs.

All rights reserved. No part of this book may be reproduced, in any form or by any means,
without permission in writing from the publisher.

Printed in the United States of America

10 9 8 7 6 5 4 3 2

ISBN 0-13-084370-9

Prentice-Hall International (UK) Limited, *London*
Prentice-Hall of Australia Pty. Limited, *Sydney*
Prentice-Hall Canada, Inc., *Toronto*
Prentice-Hall Hispanoamericana, S. A., *Mexico*
Prentice-Hall of India Private Limited, *New Delhi*
Prentice-Hall of Japan, Inc., *Tokyo*
Prentice-Hall (Singapore), Pte. Ltd., *Singapore*
Editora Prentice-Hall do Brasil, Ltda., *Rio de Janeiro*

Business Library
University of Wisconsin - Madison
2200 Grainger Hall
975 University Avenue
Madison, WI 53706-1323

The discussion so far does not touch on other key aspects of data communications, including data link control techniques for controlling the flow of data and detecting and correcting errors, and multiplexing techniques for transmission efficiency. All of these topics are explored in Part Two.

1.3 DATA COMMUNICATION NETWORKING

It is often impractical for two communicating devices to be directly, point-to-point connected. This is so for one (or both) of the following contingencies:

- The devices are very far apart. It would be inordinately expensive, for example, to string a dedicated link between two devices thousands of kilometers apart.
- There is a set of devices, each of which may require a link to many of the others at various times. Examples are all of the telephones in the world and all of the terminals and computers owned by a single organization. Except for the case of a very few devices, it is impractical to provide a dedicated wire between each pair of devices.

The solution to this problem is to attach each device to a communication network. Figure 1.3 relates this discussion to the communications model of Figure 1.1a and also suggests the two major categories into which communications networks are traditionally classified: wide area networks (WANs) and local area networks (LANs). The distinction between the two, both in terms of technology and application, has become somewhat blurred in recent years, but it remains a useful way of organizing the discussion.

Wide Area Networks

Wide area networks generally cover a large geographical area, require the crossing of public right-of-ways, and rely at least in part on circuits provided by a common carrier. Typically, a WAN consists of a number of interconnected switching nodes. A transmission from any one device is routed through these internal nodes to the specified destination device. These nodes (including the boundary nodes) are not concerned with the content of the data; rather, their purpose is to provide a switching facility that will move the data from node to node until they reach their destination.

Traditionally, WANs have implemented using one of two technologies: circuit switching and packet switching. More recently, frame relay and ATM networks have assumed major roles.

Circuit Switching

In a circuit-switching network, a dedicated communications path is established between two stations through the nodes of the network. That path is a connected sequence of physical links between nodes. On each link, a logical channel is dedicated to the connection. Data generated by the source station are transmitted along the dedicated path as rapidly as possible. At each node, incoming data are routed

10 CHAPTER 1 / INTRODUCTION

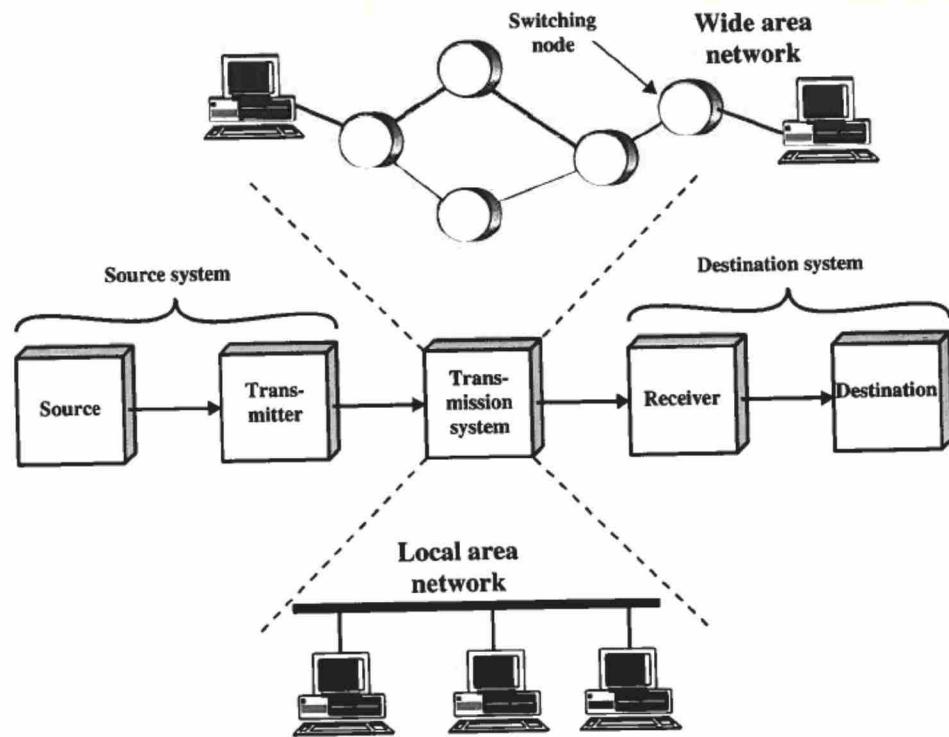


Figure 1.3 Simplified Network Models

or switched to the appropriate outgoing channel without delay. The most common example of circuit switching is the telephone network.

Packet Switching

A quite different approach is used in a packet-switching network. In this case, it is not necessary to dedicate transmission capacity along a path through the network. Rather, data are sent out in a sequence of small chunks, called packets. Each packet is passed through the network from node to node along some path leading from source to destination. At each node, the entire packet is received, stored briefly, and then transmitted to the next node. Packet-switching networks are commonly used for terminal-to-computer and computer-to-computer communications.

Frame Relay

Packet switching was developed at a time when digital long-distance transmission facilities exhibited a relatively high error rate compared to today's facilities. As a result, there is a considerable amount of overhead built into packet-switching schemes to compensate for errors. The overhead includes additional bits added to each packet to introduce redundancy and additional processing at the end stations and the intermediate switching nodes to detect and recover from errors.

With modern high-speed telecommunications systems, this overhead is unnecessary and counterproductive. It is unnecessary because the rate of errors has been dramatically lowered and any remaining errors can easily be caught in the end systems by logic that operates above the level of the packet-switching logic. It is counterproductive because the overhead involved soaks up a significant fraction of the high capacity provided by the network.

Frame relay was developed to take advantage of these high data rates and low error rates. Whereas the original packet-switching networks were designed with a data rate to the end user of about 64 kbps, frame relay networks are designed to operate efficiently at user data rates of up to 2 Mbps. The key to achieving these high data rates is to strip out most of the overhead involved with error control.

ATM

Asynchronous transfer mode (ATM), sometimes referred to as cell relay, is a culmination of all of the developments in circuit switching and packet switching over the past 25 years.

ATM can be viewed as an evolution from frame relay. The most obvious difference between frame relay and ATM is that frame relay uses variable-length packets, called frames, and ATM uses fixed-length packets, called cells. As with frame relay, ATM provides little overhead for error control, depending on the inherent reliability of the transmission system and on higher layers of logic in the end systems to catch and correct errors. By using a fixed packet length, the processing overhead is reduced even further for ATM compared to frame relay. The result is that ATM is designed to work in the range of 10s and 100s of Mbps, and in the Gbps range.

ATM can also be viewed as an evolution from circuit switching. With circuit switching, only fixed-data-rate circuits are available to the end system. ATM allows the definition of multiple virtual channels with data rates that are dynamically defined at the time the virtual channel is created. By using small, fixed-size cells, ATM is so efficient that it can offer a constant-data-rate channel even though it is using a packet switching technique. Thus, ATM extends circuit switching to allow multiple channels with the data rate on each channel dynamically set on demand.

ISDN and Broadband ISDN

Merging and evolving communications and computing technologies, coupled with increasing demands for efficient and timely collection, processing, and dissemination of information, are leading to the development of integrated systems that transmit and process all types of data. A significant outgrowth of these trends is the integrated services digital network (ISDN).

The ISDN is designed to replace existing public telecommunications networks and deliver a wide variety of services. The ISDN is defined by the standardization of user interfaces and implemented as a set of digital switches and paths supporting a broad range of traffic types and providing value-added processing services. In practice, there are multiple networks, implemented within national boundaries, but from the user's point of view, there is a single, uniformly accessible, worldwide network.

Despite the fact that ISDN has yet to achieve the universal deployment hoped for, it is already in its second generation. The first generation, sometimes referred

12 CHAPTER 1 / INTRODUCTION

to as **narrowband ISDN**, is based on the use of a 64-kbps channel as the basic unit of switching and has a circuit-switching orientation. The major technical contribution of the narrowband ISDN effort has been frame relay. The second generation, referred to as **broadband ISDN**, supports very high data rates (100s of Mbps) and has a packet-switching orientation. The major technical contribution of the broadband ISDN effort has been asynchronous transfer mode (ATM), also known as cell relay.

Local Area Networks

As with WANs, a LAN is a communications network that interconnects a variety of devices and provides a means for information exchange among those devices. There are several key distinctions between LANs and WANs:

1. The scope of the LAN is small, typically a single building or a cluster of buildings. This difference in geographic scope leads to different technical solutions, as we shall see.
2. It is usually the case that the LAN is owned by the same organization that owns the attached devices. For WANs, this is less often the case, or at least a significant fraction of the network assets are not owned. This has two implications. First, care must be taken in the choice of LAN, because there may be a substantial capital investment (compared to dial-up or leased charges for WANs) for both purchase and maintenance. Second, the network management responsibility for a LAN falls solely on the user.
3. The internal data rates of LANs are typically much greater than those of WANs.

Traditionally, LANs make use of a broadcast network approach rather than a switching approach. With a broadcast communication network, there are no intermediate switching nodes. At each station, there is a transmitter/receiver that communicates over a medium shared by other stations. A transmission from any one station is broadcast to and received by all other stations. Data are usually transmitted in packets. Because the medium is shared, only one station at a time can transmit a packet.

More recently, examples of switched LANs, especially switched Ethernet LANs, have appeared. Two other prominent examples are ATM LANs, which simply use an ATM network in a local area, and Fibre Channel. We will examine these LANs, as well as broadcast LANs, in Part Four.

1.4 PROTOCOLS AND PROTOCOL ARCHITECTURE

When computers, terminals, and/or other data processing devices exchange data, the scope of concern is much broader than the concerns we have discussed in Sections 1.2 and 1.3. Consider, for example, the transfer of a file between two computers. There must be a data path between the two computers, either directly or via a communication network. But more is needed. Typical tasks to be performed:

CHAPTER 13

LAN TECHNOLOGY

13.1 LAN Applications

13.2 LAN Architecture

13.3 BUS LANs

13.4 Ring LANs

13.5 Star LANs

13.6 Wireless LANs

13.7 Bridges

13.8 Recommended Reading and Web Sites

13.9 Problems

APPENDIX 13A The IEEE 802 Standards

- ◆ A LAN consists of a shared transmission medium and a set of hardware and software for interfacing devices to the medium and regulating the orderly access to the medium. The topologies that have been used for LANs are ring, bus, tree, and star. The bus and tree topologies are passive sections of cable to which stations are attached. A transmission of a frame by any one station can be heard by any other station. A ring LAN consists of a closed loop of repeaters that allow data to circulate around the ring. A repeater may also function as a device attachment point. Transmission is generally in the form of frames. A star LAN includes a central node to which stations are attached.
- ◆ The transmission media that are used for LANs are twisted pair, coaxial cable, optical fiber, and wireless. Both shielded and unshielded twisted pair are in use. Wireless transmission uses either infrared or microwave.
- ◆ A set of standards has been defined for LANs that specifies a range of data rates and encompasses all of the topologies and transmission media just mentioned. These standards, the IEEE 802 and fiber distributed data interface (FDDI) standards, are widely accepted, and most of the products on the market conform to one of these standards.
- ◆ In most cases, an organization will have multiple LANs that need to be interconnected. The simplest approach to meeting this requirement is the bridge.

We turn now to a discussion of local area networks (LANs). Whereas wide area networks may be public or private, LANs usually are owned by the organization that is using the network to interconnect equipment. LANs have much greater capacity than wide area networks, to carry what is generally a greater internal communications load.

A simple example of a LAN that highlights some of its characteristics is shown in Figure 1.3. All of the devices are attached to a shared transmission medium. A transmission from any one device can be received by all other devices attached to the same network. Traditional LANs have provided data rates in a range from about 1 to 20 Mbps. These data rates, though substantial, have become increasingly inadequate with the proliferation of devices, the growth in multimedia applications, and the increased use of the client/server architecture. Thus, recent efforts have focused on the development of high-speed LANs, with data rates of 100 Mbps to 1 Gbps.

This chapter begins our discussion of LANs¹ with a discussion of LAN application areas. This is followed by a description of the protocol architecture that is in common use for implementing LANs. This architecture is also the basis of standardization efforts. Our overview covers the physical, medium access control (MAC), and logical link control (LLC) levels.

¹For the sake of brevity, the chapter often uses LAN when referring to LAN and metropolitan area network (MAN) concerns. The context should clarify when only LAN or both LAN and MAN is meant.

The key technology ingredients that determine the nature of a LAN or MAN are:

- Topology
- Transmission medium
- Medium access control technique

This chapter surveys the topologies and transmission media that are most commonly used for LANs. The issue of access control is briefly raised but is covered in more detail in Chapter 14. The concept of a bridge, which plays a critical role in extending LAN coverage, is discussed in Section 13.7.

13.1 LAN APPLICATIONS

The variety of applications for LANs is wide. This section provides a brief discussion of some of the most important general application areas for these networks.

Personal Computer LANs

A common LAN configuration is one that supports personal computers. With the relatively low cost of such systems, individual managers within organizations often independently procure personal computers for departmental applications, such as spreadsheet and project management tools, and Internet access.

But a collection of department-level processors will not meet all of an organization's needs; central processing facilities are still required. Some programs, such as econometric forecasting models, are too big to run on a small computer. Corporate-wide data files, such as accounting and payroll, require a centralized facility but should be accessible to a number of users. In addition, there are other kinds of files that, although specialized, must be shared by a number of users. Further, there are sound reasons for connecting individual intelligent workstations not only to a central facility but to each other as well. Members of a project or organization team need to share work and information. By far the most efficient way to do so is digitally.

Certain expensive resources, such as a disk or a laser printer, can be shared by all users of the departmental LAN. In addition, the network can tie into larger corporate network facilities. For example, the corporation may have a building-wide LAN and a wide area private network. A communications server can provide controlled access to these resources.

LANs for the support of personal computers and workstations have become nearly universal in organizations of all sizes. Even those sites that still depend heavily on the mainframe have transferred much of the processing load to networks of personal computers. Perhaps the prime example of the way in which personal computers are being used is to implement client/server applications.

For personal computer networks, a key requirement is low cost. In particular, the cost of attachment to the network must be significantly less than the cost of the attached device. Thus, for the ordinary personal computer, an attachment cost in the hundreds of dollars is desirable. For more expensive, high-performance workstations, higher attachment costs can be tolerated. In any case, this suggests that the

Appendix H

"I just don't see how anyone can run their network without it."

—Teré Parnell, Executive Technology Editor, *LAN Times*



FEATURING FULLY SEARCHABLE,
HYPERLINKED CD-ROM

LAN TIMES

OSBORNE

Encyclopedia of Networking

• ELECTRONIC EDITION

TOM SHELDON

Certified Network Engineer and Author of the Best-selling
LAN Times Encyclopedia of Networking

- Thousands of Networking Terms and Concepts Thoroughly Explained

- Every Major Networking System Covered! UNIX, Windows NT, Novell NetWare, IBM, the Internet & More

- The Ideal Reference for Everyone Involved with Networks and the Internet



Encyclopedia of Networking, Electronic Edition

Tom Sheldon

Osborne McGraw-Hill

Berkeley New York St. Louis San Francisco
Auckland Bogotá Hamburg London Madrid
Mexico City Milan Montreal New Delhi Panama City
Paris São Paulo Singapore Sydney
Tokyo Toronto

Osborne/McGraw-Hill
2600 Tenth Street
Berkeley, California 94710
U.S.A.

For information on translations or book distributors outside the U.S.A., or to arrange bulk purchase discounts for sales promotions, premiums, or fund-raisers, please contact Osborne/McGraw-Hill at the above address.

Encyclopedia of Networking, Electronic Edition

Copyright © 1998 by The McGraw-Hill Companies. All rights reserved. Printed in the United States of America. Except as permitted under the Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written permission of the publisher, with the exception that the program listings may be entered, stored, and executed in a computer system, but they may not be reproduced for publication.

567890 DOC DOC 9019876543210

ISBN 0-07-882333-1 PPBK

ISBN 0-07-882350-1 HRDBK

Publisher

Brandon A. Nordin

Proofreader

Karen Mead

Editor-in-Chief

Scott Rogers

Indexer

Dan Logan

Acquisitions Editor

Wendy Rinaldi

Computer Designer

Jani Beckwith

Project Editor

Emily Rader

Illustrators

Sue Albert

Leslee Bassin

Arlette Crosland

Lance Ravella

Editorial Assistant

Ann Sellers

Cover Design

Regan Honda

Technical Editor

Terè Parnell

Copy Editor

Dennis Weaver

Information has been obtained by Osborne/McGraw-Hill from sources believed to be reliable. However, because of the possibility of human or mechanical error by our sources, Osborne/McGraw-Hill, or others, Osborne/McGraw-Hill does not guarantee the accuracy, adequacy, or completeness of any information and is not responsible for any errors or omissions or the results obtained from use of such information.

L2F (Layer 2 Forwarding)

See Virtual Dial-up Services.

L2TP (Layer 2 Tunnelling Protocol)

See Virtual Dial-up Services.

Label Switching

See IP over ATM; IP Switching; MPLS (Multiprotocol Label Switching); and T1g Switching.

LAN (Local Area Network)

A LAN is a shared communication system to which many computers are attached. A LAN, as its name implies, is limited to a local area. This has to do more with the electrical characteristics of the medium than the fact that many early LANs were designed for departments, although the latter accurately describes a LAN as well.

LANs began to appear in the early 1970s. They grew from earlier point-to-point connections where a single wire connected two systems. Often the wire was quite long. Why not let multiple computers share this same cable? This required an arbitration mechanism to ensure that only one computer transmitted at once on the cable.

Arbitration methods are called *medium access controls*. Some methods have each workstation determine whether the cable is in use. Other methods use a central controller that gives each station access in turn. See "MAC (Medium Access Control) and Medium Access Control Methods" for more information on access methods.

LANs have different topologies, the most common being the *linear bus* and the *star configuration*. In the former, a cable snakes through a building from one workstation to another. In the star configuration, each workstation is connected to a central hub with its own cable. Each has its advantages and disadvantages. Interestingly, the most popular network, Ethernet, can take advantage of both topologies. Refer to "Topology" for more details.

A LAN is a connectionless networking scheme, meaning that once a workstation is ready to transmit and has access to the shared medium, it simply puts the packets on the network and hopes that the recipient receives them. There is no connection setup phase in this scheme. See "Connection-Oriented and Connectionless Services" for more details.

Data is packaged into *frames* for transmission on the LAN. At the hardware level, each frame is transmitted as a bit stream on the wire. Even though all the computers on the network listen to the transmission, only the designated recipient actually receives the frame. A frame is usually addressed for a single computer, although a *multicast address* can be used to transmit to all workstations on the LAN. Higher-layer protocols such as IP and IPX package data into datagrams. Datagrams are in turn

divided up and put into frames for transmission on a particular LAN. See "Datagrams and Datagram Services" and "Framing in Data Transmissions" for more details.

LAN Distance and Size Limitations

One of the reasons why LANs are considered "local" is because there are practical limitations to the distance of a shared medium and the number of workstations you can connect to it. For example, if you tried to build a single LAN for an entire organization, there might be so many workstations attempting to access the cable at the same time that no real work would get done.

The electrical characteristics of the cable also dictate LAN limitations. Network designers must find a balance among the type of cable used, the transmission rates, signal loss over distance, and the signal emanation. All of these factors must stay within physical bounds and restrictions specified by various standards and government bodies. For example, coaxial cable allows higher transmission rates over longer distances, but twisted-pair wire is inexpensive, easy to install, and supports a hierarchical wiring scheme.

Delay is another factor. On Ethernet networks, workstations on either end of a long cable may not even detect that they are transmitting at the same time, thus causing a collision that results in corrupted data. You can use the following devices to extend a LAN or improve its performance:

- **Repeaters** Extends the limitations of Ethernet cable by boosting the signal. See "Repeater" for details.
- **Bridges** Provides repeater functions along with selective filtering of traffic to reduce congestion and contention. See "Bridges and Bridging" for details.
- **Switching** Provides an overall improvement in LAN throughput and design as described under "Switched Networks."
- **Routers** Provide a way to connect multiple LANs together to create internetworks. See "Internetworking" and "IP (Internet Protocol)" for more details.

RELATED ENTRIES Bridges and Bridging; Broadcast Networking; Connection Technologies; Data Communication Concepts; Datagram, and Datagram Services; Data Link Protocols; Ethernet; Framing in Data Transmissions; Internetworking; LAN Emulation; MAC (Medium Access Control); Medium Access Control Methods; Network Concepts; Network Design and Construction; Network Operating Systems; Packet; Protocol Concepts; Repeater; Switched Networks; Token Ring Network; Topology; VLAN (Virtual LAN); and Wireless Communications

LAN Drivers

A LAN driver is a workstation or server software module that provides an interface between a NIC (network interface card) and the upper-layer protocol software running in the computer. The driver is designed for a specific NIC. Drivers are usually

Protocols specify a set of rules and procedures that define how communication takes place at different levels of operation. The lowest layers define physical connections, such as the cable type, access method, and topology, and how data is sent over the network. Further up are protocols that establish connections and maintain communication sessions between systems, and still further up are protocols that provide network interfacing for applications.

As mentioned, the OSI model has become the model to which all other network architectures and protocols are compared. The purpose of the OSI model is to coordinate communication standards between vendors. Refer to "OSI (Open Systems Interconnection) Model" for additional information.

RELATED ENTRIES Data Communication Concepts; Network Design and Construction; OSI (Open Systems Interconnection) Model; and Protocol Concepts

Network Computer Devices

See NC (Network Computer) Devices.

Network Concepts

A network is a communication system that allows users to access resources on other computers and exchange messages with other users. It allows users to share resources on their own systems with other network users and to access information on centrally located systems or systems that are located at remote offices. It may provide connections to the Internet or the networks of other organizations. Network connections allow users to operate from their home or on the road.

The Scope of Networks

A network is a data communication system that links two or more computers and peripheral devices. As shown in Figure N-10, the network consists of a cable that attaches to NICs (network interface cards) in each of the devices. Figure N-11 illustrates the logical configuration of a network communication system. Users interact with network-enabled software applications to make a network request (such as to get a file or print on a network printer). The application communicates with the network software and the network software interacts with the network hardware. The network hardware is responsible for transmitting information to other devices attached to the network.

LAN (Local Area Network)

A LAN is a network that is located in a relatively small area, such as a department or building. Technically, a LAN consists of a shared medium to which workstations attach and communicate with one another using broadcast methods. With broadcasting, any device on the LAN can transmit a message that all other devices on the LAN can

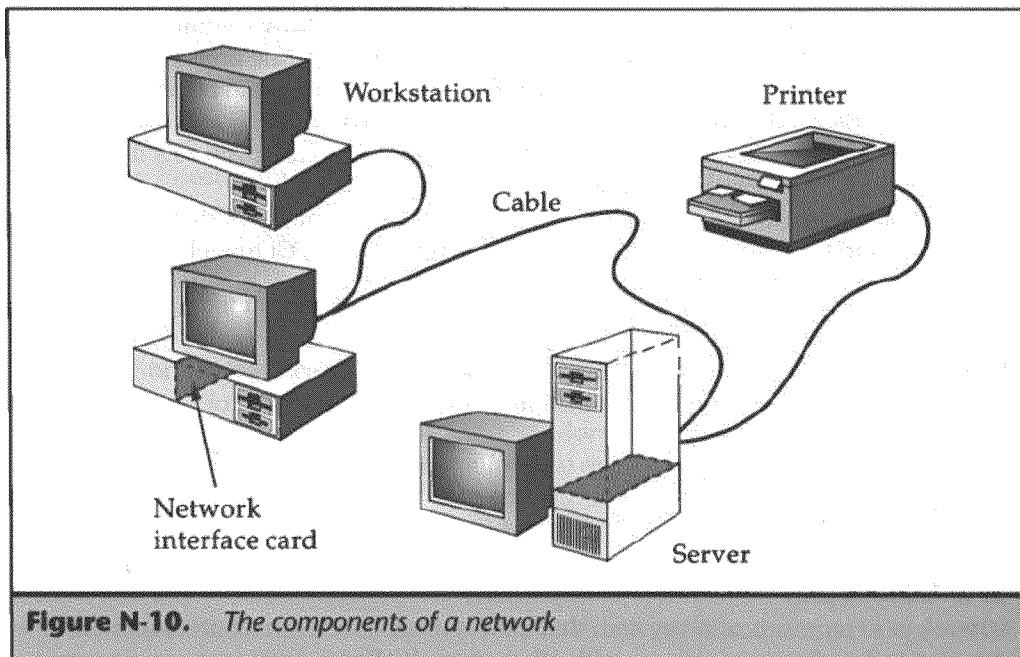


Figure N-10. The components of a network

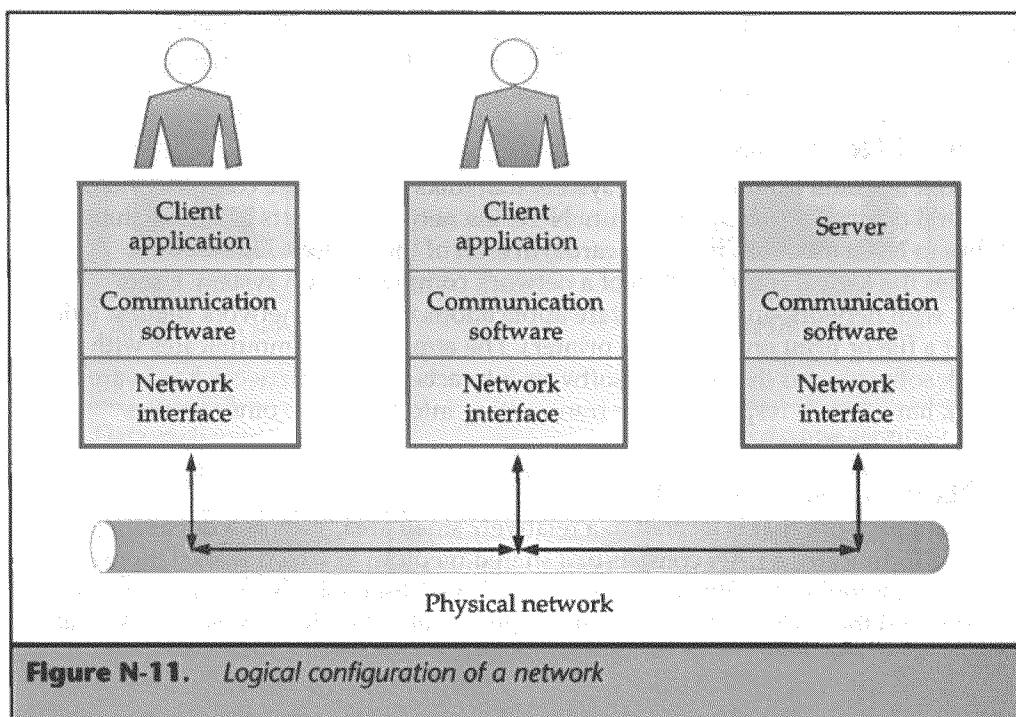


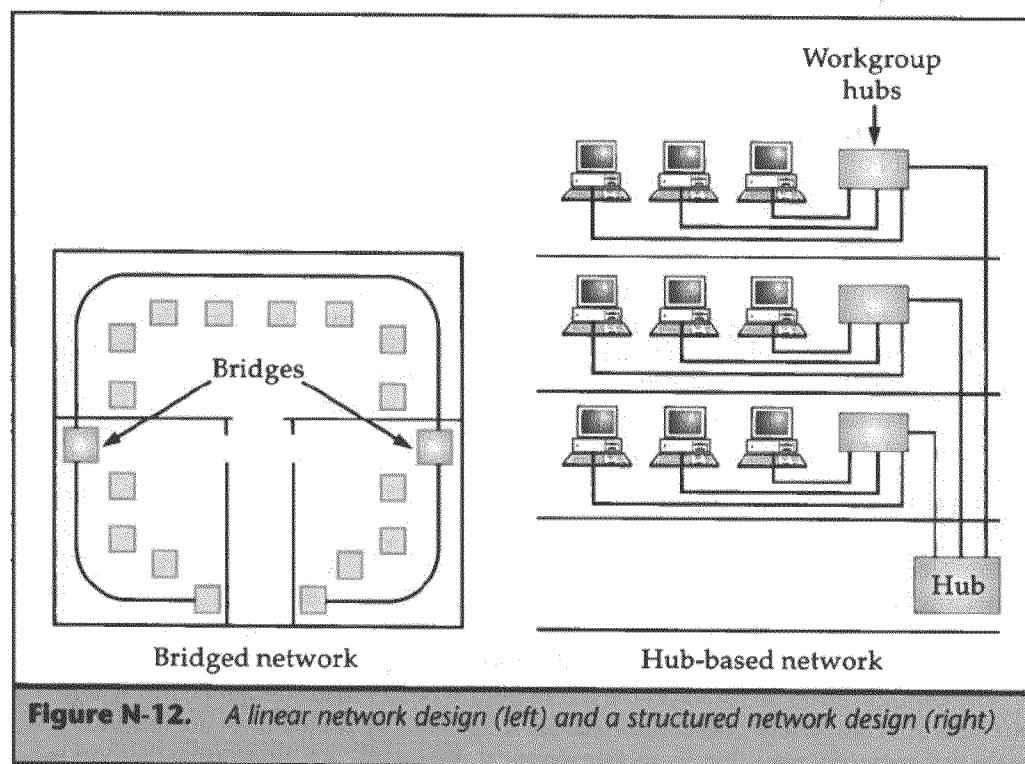
Figure N-11. Logical configuration of a network

listen to. The device to which the message is addressed actually receives the message. See "LAN (Local Area Network)" for more details.

Figure N-12 illustrates two ways to build a LAN. On the left, a relatively small LAN is built by running cable in a daisy-chain fashion from one department to the next. Each LAN segment is joined with a bridge. A bridge extends a LAN to create a much larger broadcast domain, but the bridge filters each individual segment's broadcasts by dropping frames that are not addressed to devices on connected segments. On the right, several LANs are interconnected at a centrally located hub device that handles the delivery of all inter-LAN traffic. See "Bridges and Bridging" and "Hubs/Concentrators/MAUs" for more information.

The model on the right in Figure N-13 implements a structured wiring system that is hierarchical in nature. Cables branch from a central internetwork hub to departmental hubs. This system of interconnecting cables and hub is often referred to as the backbone network. See "Backbone Networks" for more information.

The two most popular LAN technologies are Ethernet and token ring. See "Ethernet" and "Token Ring Network" for more details.



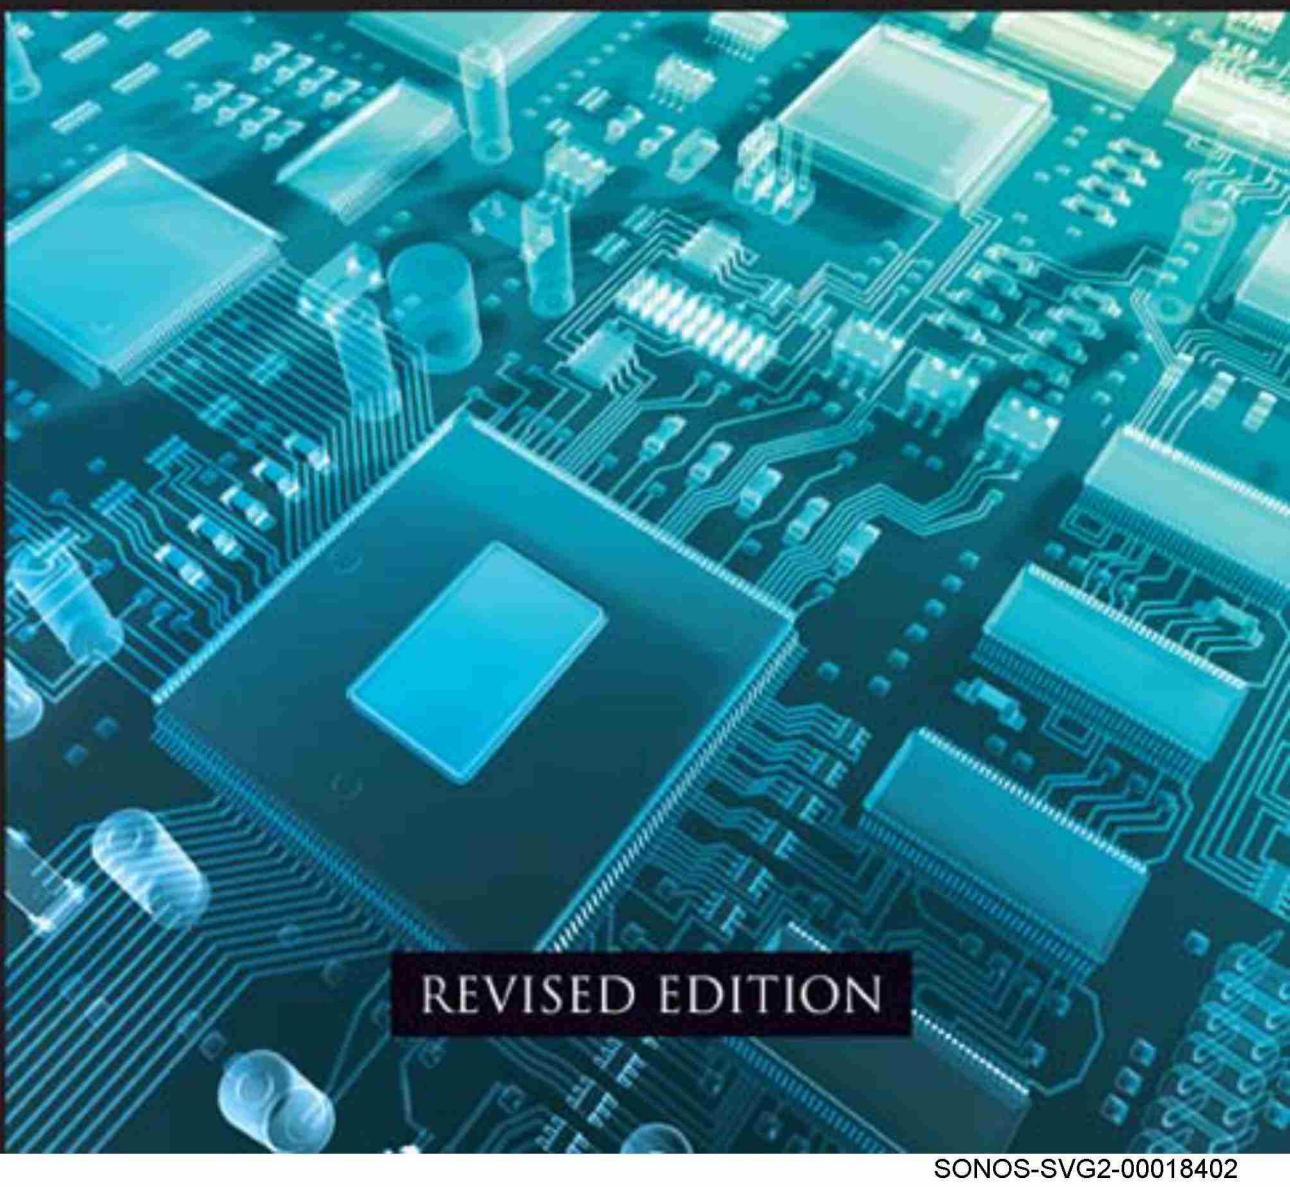
Appendix I

FACTS ON FILE SCIENCE LIBRARY



ENCYCLOPEDIA OF COMPUTER SCIENCE AND TECHNOLOGY

HARRY HENDERSON



REVISED EDITION

ENCYCLOPEDIA OF COMPUTER SCIENCE AND TECHNOLOGY

REVISED EDITION

HARRY HENDERSON

 **Facts On File**
An imprint of Infobase Publishing

*In memory of my brother,
Bruce Henderson,
who gave me my first opportunity to explore
personal computing almost 30 years ago.*

ENCYCLOPEDIA OF COMPUTER SCIENCE AND TECHNOLOGY, Revised Edition

Copyright © 2009, 2004, 2003 by Harry Henderson

All rights reserved. No part of this book may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage or retrieval systems, without permission in writing from the publisher.

For information contact:

Facts On File, Inc.
An imprint of Infobase Publishing
132 West 31st Street
New York NY 10001

Library of Congress Cataloging-in-Publication Data

Henderson, Harry, 1951–
Encyclopedia of computer science and technology / Harry Henderson.—Rev. ed.
p. cm.

Includes bibliographical references and index.

ISBN-13: 978-0-8160-6382-6

ISBN-10: 0-8160-6382-6

1. Computer science—Encyclopedias. 2. Computers—Encyclopedias. I. Title.
QA76.15.H43 2008

004.03—dc22

2008029156

Facts On File books are available at special discounts when purchased in bulk quantities for businesses, associations, institutions, or sales promotions. Please call our Special Sales Department in New York at (212) 967-8800 or (800) 322-8755.

You can find Facts On File on the World Wide Web at <http://www.factsonfile.com>

Text design by Erika K. Arroyo
Cover design by Salvatore Luongo
Illustrations by Sholto Ainslie
Photo research by Tobi Zausner, Ph.D.

Printed in the United States of America

VB Hermitage 10 9 8 7 6 5 4 3 2 1

This book is printed on acid-free paper and contains
30 percent postconsumer recycled content.

Petzold, Charles. *Programming Windows: the Definitive Guide to the Win32 API*. 5th ed. Redmond, Wash.: Microsoft Press, 1999. "Windows API Guide." Available online. URL: <http://www.vbapi.com/>. Accessed April 12, 2007.

application service provider (ASP)

Traditionally, software applications such as office suites are sold as packages that are installed and reside on the user's computer. Starting in the mid-1990s, however, the idea of offering users access to software from a central repository attracted considerable interest. An application service provider (ASP) essentially rents access to software.

Renting software rather than purchasing it outright has several advantages. Since the software resides on the provider's server, there is no need to update numerous desktop installations every time a new version of the software (or a "patch" to fix some problem) is released. The need to ship physical CDs or DVDs is also eliminated, as is the risk of software piracy (unauthorized copying). Users may be able to more efficiently budget their software expenses, since they will not have to come up with large periodic expenses for upgrades. The software provider, in turn, also receives a steady income stream rather than "surges" around the time of each new software release.

For traditional software manufacturers, the main concern is determining whether the revenue obtained by providing its software as a service (directly or through a third party) is greater than what would have been obtained by selling the software to the same market. (It is also possible to take a hybrid approach, where software is still sold, but users are offered additional features online. Microsoft has experimented with this approach with its Microsoft Office Live and other products.)

Renting software also has potential disadvantages. The user is dependent on the reliability of the provider's servers and networking facilities. If the provider's service is down, then the user's work flow and even access to critical data may be interrupted. Further, sensitive data that resides on a provider's system may be at risk from hackers or industrial spies. Finally, the user may not have as much control over the deployment and integration of software as would be provided by outright purchase.

The ASP market was a hot topic in the late 1990s, and some pundits predicted that the ASP model would eventually supplant the traditional retail channel for mainstream software. This did not happen, and more than a thousand ASPs were among the casualties of the "dot-com crash" of the early 2000s. However, ASP activity has been steadier if less spectacular in niche markets, where it offers more economical access to expensive specialized software for applications such as customer relationship management, supply chain management, and e-commerce related services—for example, Salesforce.com. The growing importance of such "software as a service" business models can be seen in recent offerings from traditional software companies such as SAS. By 2004, worldwide spending for "on demand" software had exceeded \$4 billion, and Gartner Research has predicted that in the second half of the decade about

a third of all software will be obtained as a service rather than purchased.

WEB-BASED APPLICATIONS AND FREE SOFTWARE

By that time a new type of application service provider had become increasingly important. Rather than seeking to gain revenue by selling online access to software, this new kind of ASP provides the software for free. A striking example is Google Pack, a free software suite offered by the search giant (see GOOGLE). Google Pack includes a variety of applications, including a photo organizer and search and mapping tools developed by Google, as well as third-party programs such as the Mozilla Firefox Web browser, Real-Player media player, the Skype Internet phone service (see VOIP), and antivirus and antispyware programs. The software is integrated into the user's Windows desktop, providing fast index and retrieval of files from the hard drive. (Critics have raised concerns about the potential violation of privacy or misuse of data, especially with regard to a "share across computers" feature that stores data about user files on Google's servers.) America Online has also begun to provide free access to software that was formerly available only to paid subscribers.

This use of free software as a way to attract users to advertising-based sites and services could pose a major threat to companies such as Microsoft that rely on software as their main source of revenue. In 2006 Google unveiled a Google Docs & Spreadsheets, a program that allows users to create and share word-processing documents and spreadsheets over the Web. Such offerings, together with free open-source software such as Open Office.org, may force traditional software companies to find a new model for their own offerings.

Microsoft in turn has launched Office Live, a service designed to provide small offices with a Web presence and productivity tools. The free "basic" level of the service is advertising supported, and expanded versions are available for a modest monthly fee. The program also has features that are integrated with Office 2007, thus suggesting an attempt to use free or low-cost online services to add value to the existing stand-alone product line.

By 2008 the term *cloud computing* had become a popular way to describe software provided from a central Internet site that could be accessed by the user through any form of computer and connection. An advantage touted for this approach is that the user need not be concerned with where data is stored or the need to make backups, which are handled seamlessly.

Further Reading

Chen, Anne. "Office Live Makes Online Presence Known." *eWeek*, November 2, 2006. Available online. URL: <http://www.eweek.com/article2/0,1759,2050580,00.asp>. Accessed May 22, 2007.

Focacci, Luisa, Robert J. Mockler, and Marc E. Gartenfeld. *Application Service Providers in Business*. New York: Haworth, 2005.

Garretson, Rob. "The ASP Reincarnation: The Application Service Provider Name Dies Out, but the Concept Lives on among Second-Generation Companies Offering Software as a service." *Network World*, August 29, 2005. Available online.

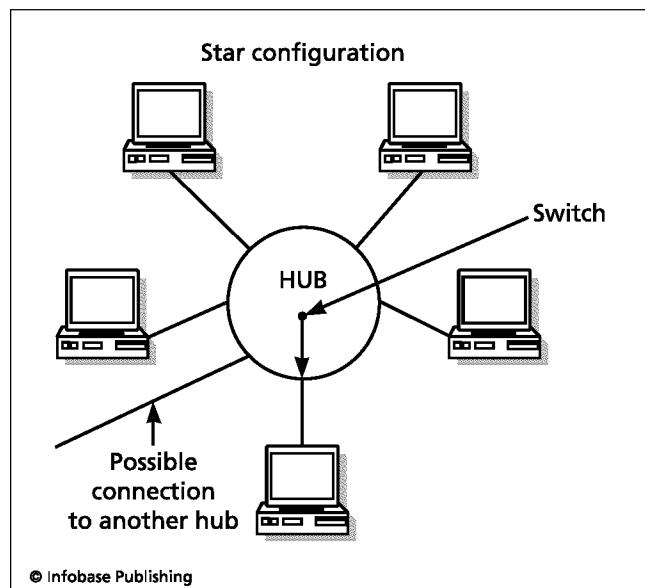
edu/mc/LispNotes/RecursiveListProcessingAlgorithmsIn-Lisp.pdf. Accessed August 13, 2007.
 "Linked List Basics [In C/C++]" Available online. URL: <http://cslibrary.stanford.edu/103/LinkedListBasics.pdf>. Accessed August 13, 2007.
 Seibel, Peter. "They Called It LISP for a Reason: List Processing" Available online. URL: <http://www.gigamonkeys.com/book/they-called-it-lisp-for-a-reason-list-processing.html>. Accessed August 13, 2007.

local area network (LAN)

Starting in the 1980s, many organizations sought to connect their employees' desktop computers so they could share central databases, share or back up files, communicate via e-mail, and collaborate on projects. A system that links computers within a single office or home, or a larger area such as a building or campus, is called a local area network (LAN). (Larger networks linking branches of an organization throughout the country or world are called wide area networks, or WANs. See NETWORK.)

HARDWARE ARCHITECTURE

There are two basic ways to connect computers in a LAN. The first, called Ethernet, was developed by a project at the Xerox Palo Alto Research Center (PARC) led by Robert Metcalfe. Ethernet uses a single cable line called a bus to which all participating computers are connected. Each data packet is received by all computers, but processed only by the one it is addressed to. Before sending a packet, a computer first checks to make sure the line is free. Sometimes, due to the time delay before a packet is received by all computers, another computer may think the line is free and start trans-

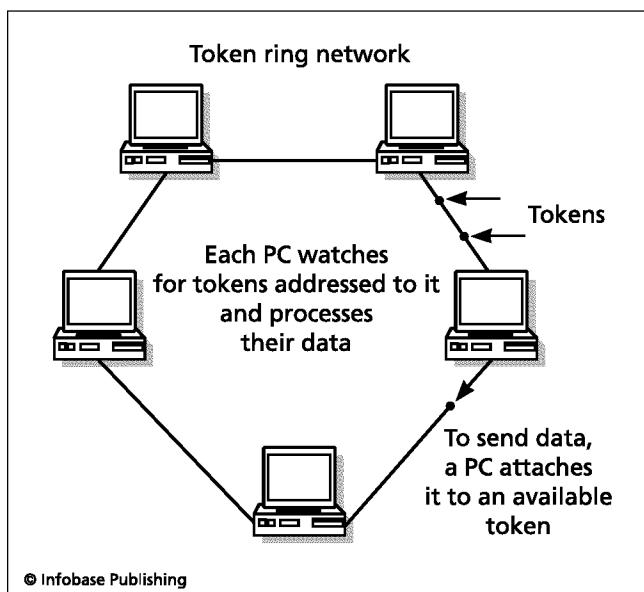


The Star network configuration uses a central hub to which each PC is attached. To extend the network (such as into other offices), the hubs can be connected to one another so they function as switches. When a token arrives that is addressed to one of its PCs, the hub will route it to the appropriate machine.

mitting. The resulting *collision* is resolved by having both computers stop and wait varying times before resending.

Because connecting all computers to a single bus line is impractical in larger installations, Ethernet networks are frequently extended to multiple offices by connecting a bus in each office to a switch, creating a subnetwork or segment (this is sometimes called a *star topology*). The switches are then connected to a main bus. Packets are first routed to the switch for the segment containing the destination computer. The switch then dispatches the packet to the destination computer. Another advantage of this *switched Ethernet* system is that more-expensive, high-bandwidth cable can be used to connect the switches to move the packets more quickly over greater distances, while less-expensive cabling can be used to connect each computer to its local switch.

An alternative way to arrange a LAN is called *token ring*. Instead of the computers being connected to a bus that ends in a terminator, they are connected in a circle where the last computer is connected to the first. Interference is prevented by using a special packet called the token. Like the use of a "talking stick" in a tribal council, only the computer holding the token can transmit at a given time. After transmitting, the computer puts the token back into circulation so it can be grabbed by the next computer that wants to send data.



A Token Ring network connects the machines in a "chain" around which messages called tokens travel. Any PC can "grab" a passing token and attach data and the address of another PC to it. Each PC in turn watches for tokens that are addressed to it.

LAN SOFTWARE

Naturally there must be software to manage the transmission and reception of data packets. The structure of a packet (sometimes called a *frame*) has been standardized with a preamble, source and destination addresses, the data itself,

Appendix J

SIXTH EDITION

LOCAL & METROPOLITAN AREA NETWORKS

Gigabit
Ethernet

Switched
Ethernet

Wireless
LANs

LAN
QoS

Storage
Area
Networks

WILLIAM STALLINGS

LOCAL AND METROPOLITAN AREA NETWORKS

SIXTH EDITION

William Stallings



Prentice Hall
Upper Saddle River, New Jersey 07458

Library of Congress Cataloging-in-Publication Data

Stallings, William.

Local and Metropolitan Area Networks / William Stallings.

p. cm.

Includes bibliographical references.

ISBN 0-13-012939-0

1. Computer networks. 2. LAN/MAN. I. Title.

TK5105.59.S725 1999

005.8—dc21

067829

CIP

Editor-in-chief: Marcia Horton

Publisher: Alan Apt

Project manager: Ana Terry

Executive managing editor: Vince O'Brien

Managing editor: David A. George

Vice president of production and manufacturing: David W. Riccardi

Editorial assistant: Toni Holm

Editorial/production supervision: Rose Kernan

Art director: Heather Scott

Manufacturing buyer: Pat Brown

Marketing manager: Jennie Burger



© 2000, 1997 by Prentice-Hall, Inc.
Upper Saddle River, New Jersey 07458

All right reserved. No part of this book may be reproduced, in any form or by any means, without permission in writing from the publisher.

The author and publisher of this book have used their best efforts in preparing this book. These efforts include the development, research, and testing of the theories and programs to determine their effectiveness. The author and publisher make no warranty of any kind, expressed or implied, with regard to these programs or the documentation contained in this book. The author and publisher shall not be liable in any event for incidental or consequential damages in connection with, or arising out of, the furnishing, performance, or use of these programs.

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

ISBN 0-13-012939-0

Prentice-Hall International (UK) Limited, *London*

Prentice-Hall of Australia Pty. Limited, *Sydney*

Prentice-Hall Canada, Inc., *Toronto*

Prentice-Hall Hispanoamericana, S. A., *Mexico*

Prentice-Hall of India Private Limited, *New Delhi*

Prentice-Hall of Japan, Inc., *Tokyo*

Pearson Education Asia Pte. Ltd., *Singapore*

Editora Prentice-Hall do Brasil, Ltda., *Rio de Janeiro*

CHAPTER 1

INTRODUCTION

1.1 The Need for Local Networks

1.2 LANs, MANs, and WANs

- Wide Area Network
- Local Area Network
- Metropolitan Area Networks

1.3 Applications of LANs and MANs

- Personal Computer Local Networks
- Backend Networks and Storage Area Networks
- High-Speed Office Networks
- Backbone Local Networks
- Factory Local Networks

1.4 Local Network Architecture

- Information Distribution
- Tiered LANs
- Evolution Scenario

1.5 LANs, WANs, and the Internet

1.6 Recommended Reading

1.7 Problems

Appendix 1A Internet and Web Resources

- Web Site for This Book
- Other Web Sites
- USENET Newsgroups

2 CHAPTER 1 / INTRODUCTION

The local area network (LAN) has come to play a central role in information distribution and office functioning within businesses and other organizations. The major factors driving the widespread use of LANs have been the proliferation of personal computers, workstations, and servers, coupled with the increasing reliance on the client/server computing model.

With the dropping price of LAN hardware and software, LANs have become more numerous and larger, and they have taken on more and more functions within the organization. The upshot is that the LAN, once installed, quickly becomes almost as essential as the telephone system. At the same time, there is a proliferation of LAN types and options and a need to interconnect a number of LANs at the same site and with LANs at other sites. This has led to the development of LANs of higher and higher data rates and the relatively recent introduction of the metropolitan area network (MAN).

Before defining the terms LAN and MAN, it is useful to look at the trends responsible for the importance of these networks, which we do in the first section. Next we contrast the differences among LANs, MANs, and wide area networks (WANs). This is followed by a discussion of key application areas for LANs and MANs. This chapter also provides pointers to Internet resources relating to LANs and MANs.

1.1 THE NEED FOR LOCAL NETWORKS

Perhaps the driving force behind the widespread use of LANs and MANs is the dramatic and continuing decrease in computer hardware costs, accompanied by an increase in computer hardware capability. Year by year, the cost of computer systems continues to drop dramatically while the performance and capacity of those systems continue to rise equally dramatically. At a local warehouse club, you can pick up a personal computer for less than \$1000 that packs the wallop of an IBM mainframe from 10 years ago. Inside that personal computer, including the microprocessor and memory and other chips, you get over 100 million transistors. You can't buy 100 million of anything else for so little. That many sheets of toilet paper would run more than \$100,000.

Thus we have virtually "free" computer power. And this ongoing technological revolution has enabled the development of applications of astounding complexity and power. For example, desktop applications that require the great power of today's microprocessor-based systems include

- Image processing
- Speech recognition
- Videoconferencing
- Multimedia authoring
- Voice and video annotation of files

Workstation systems now support highly sophisticated engineering and scientific applications, as well as simulation systems, and the ability to apply workgroup

principles to image and video applications. In addition, businesses are relying on increasingly powerful servers to handle transaction and database processing and to support massive client/server networks that have replaced the huge mainframe computer centers of yesteryear.

All of these factors lead to an increased number of systems, with increased power, at a single site: office building, factory, operations center, and so on. At the same time, there is an absolute requirement to interconnect these systems to

- Share and exchange data among systems
- Share expensive resources

The need to share data is a compelling reason for interconnection. Individual users of computer resources do not work in isolation. They need facilities to exchange messages with other users, to access data from several sources in the preparation of a document or for an analysis, and to share project-related information with other members of a workgroup.

The need to share expensive resources is another driving factor in the development of networks. The cost of processor hardware has dropped far more rapidly than the cost of mass storage devices, video equipment, printers, and other peripheral devices. The result is a need to share these expensive devices among a number of users to justify the cost of the equipment. This sharing requires some sort of client/server architecture operating over a network that interconnects users and resources.

1.2 LANs, MANs, AND WANs

LANs, MANs, and WANs are all examples of communications networks. A communications network is a facility that interconnects a number of devices and provides a means for transmitting data from one attached device to another.

There are a number of ways of classifying communications networks. One way is in terms of the technology used: specifically, in terms of topology and transmission medium. That approach is explored in Chapter 4. Perhaps the most commonly used means of classification is on the basis of geographical scope. Traditionally, networks have been classified as either LANs or WANs. A term that is sometimes used is the MAN.

Figure 1.1 illustrates these categories, plus some special cases. By way of contrast, the typical range of parameters for a multiple-processor computer is also depicted.

Wide Area Network

WANs have traditionally been considered to be those that cover a large geographical area, require the crossing of public right-of-ways, and rely at least in part on circuits provided by a common carrier. Typically, a WAN consists of a number of interconnected switching nodes. A transmission from any one device is routed through these internal nodes to the specified destination device.

Traditionally, WANs have provided only relatively modest capacity to subscribers. For data attachment, either to a data network or to a telephone network

4 CHAPTER 1 / INTRODUCTION

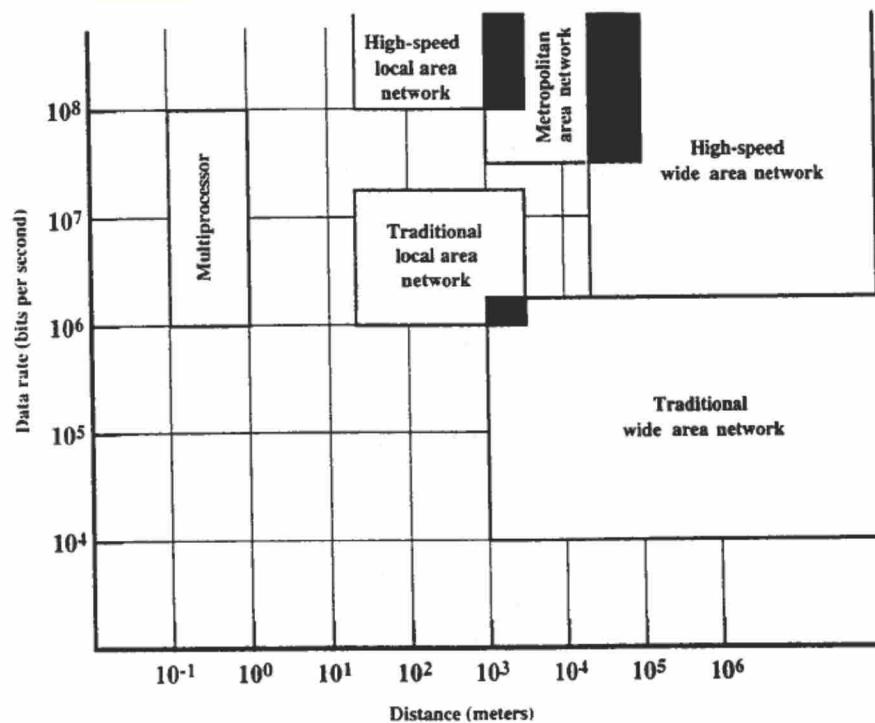


Figure 1.1 Comparison of Multiprocessor Systems, LANs, MANs, and WANs

by means of a modem, data rates of 9600 bps or even less have been common. Business subscribers have been able to obtain higher rates, with a service known as T1, which operates at 1.544 Mbps, being common. The most important recent development in WANs in this range of performance has been the development of the integrated services digital network (ISDN), which provides circuit-switching and packet-switching services at rates up to 1.544 Mbps (2.048 Mbps in Europe).

The continuing development of practical optical fiber facilities has led to the standardization of much higher data rates for WANs, and these services are becoming more widely available. These high-speed WANs provide user connections in the 10s and 100s of Mbps, using transmission techniques known as frame relay and asynchronous transfer mode (ATM).

Local Area Network

As with WANs, a LAN is a communications network that interconnects a variety of devices and provides a means for information exchange among those devices. There are several key distinctions between LANs and WANs:

1. The scope of the LAN is small, typically a single building or a cluster of buildings. This difference in geographic scope leads to different technical solutions, as we shall see.

2. It is usually the case that the LAN is owned by the same organization that owns the attached devices. For WANs, this is less often the case, or at least a significant fraction of the network assets are not owned. This has two implications. First, care must be taken in the choice of LAN, since there may be a substantial capital investment (compared to dial-up or leased charges for WANs) for both purchase and maintenance. Second, the network management responsibility for a LAN falls solely on the user.
3. The internal data rates of LANs are typically much greater than those of WANs.

LANs have been the focus of a standardization effort by the IEEE 802 committee, and it is useful to review the IEEE definition of a LAN (Table 1.1).

A Simple LAN

A simple example of a LAN that highlights some of its characteristics is shown in Figure 1.2. All of the devices are attached to a shared transmission medium. A transmission from any one device can be received by all other devices attached to the same network.

What is not apparent in Figure 1.2 is that each device attaches to the LAN through a hardware/software module that handles the transmission and medium access functions associated with the LAN. Typically, this module is implemented as a physically distinct network interface card (NIC) in each attached device. The NIC contains the logic for accessing the LAN and for sending and receiving blocks of data on the LAN.

An important function of the NIC is that it uses a buffered transmission technique to accommodate the difference in the data rate between the LAN medium

Table 1.1 Definitions of LANs and MANs*

The LANs described herein are distinguished from other types of data networks in that they are optimized for a moderate size geographic area such as a single office building, a warehouse, or a campus. The IEEE 802 LAN is a shared medium peer-to-peer communications network that broadcasts information for all stations to receive. As a consequence, it does not inherently provide privacy. The LAN enables stations to communicate directly using a common physical medium on a point-to-point basis without any intermediate switching node being required. There is always need for an access sublayer in order to arbitrate the access to the shared medium. The network is generally owned, used, and operated by a single organization. This is in contrast to Wide Area Networks (WANs) that interconnect communication facilities in different parts of a country or are used as a public utility. These LANs are also different from networks, such as backplane buses, that are optimized for the interconnection of devices on a desk top or components within a single piece of equipment.

A MAN is optimized for a larger geographical area than a LAN, ranging from several blocks of buildings to entire cities. As with local networks, MANs can also depend on communications channels of moderate-to-high data rates. Error rates and delay may be slightly higher than might be obtained on a LAN. A MAN might be owned and operated by a single organization, but usually will be used by many individuals and organizations. MANs might also be owned and operated as public utilities. They will often provide means for internetworking of local networks. Although not a requirement for all LANs, the capability to perform local networking of integrated voice and data (IVD) devices is considered an optional function for a LAN. Likewise, such capabilities in a network covering a metropolitan area are optional functions of a MAN.

* From IEEE 802 Standard, *Local and Metropolitan Area Networks: Overview and Architecture*, 1990.

6 CHAPTER 1 / INTRODUCTION

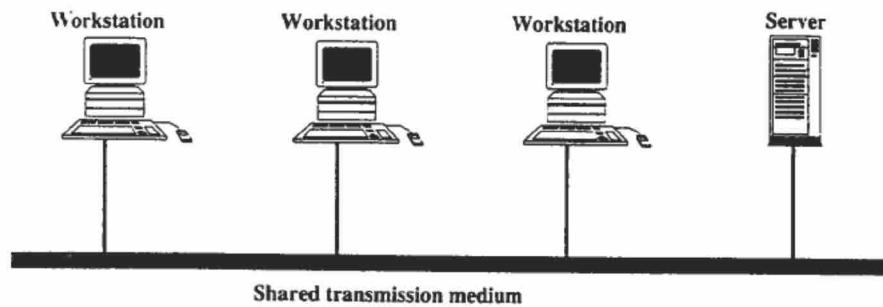


Figure 1.2 Simple Local Area Network

and the NIC-processor link, as illustrated in Figure 1.3. The NIC captures transmissions intended for the attached device, which arrive at the data rate of the LAN, which may be, for example, 10 Mbps. When a block of data is captured, it is stored temporarily in an input buffer. It is then delivered to the host processor, often over some sort of backplane bus, at the data rate of that bus. This data rate is typically different from the LAN data rate. For example, it may be 50 or 100 Mbps. Thus the NIC acts as an adapter between the data rate on the host system bus and the data rate on the LAN.

High-Speed LANs

Traditional LANs have provided data rates in a range from about 1 to 20 Mbps. These data rates, though substantial, have become increasingly inadequate

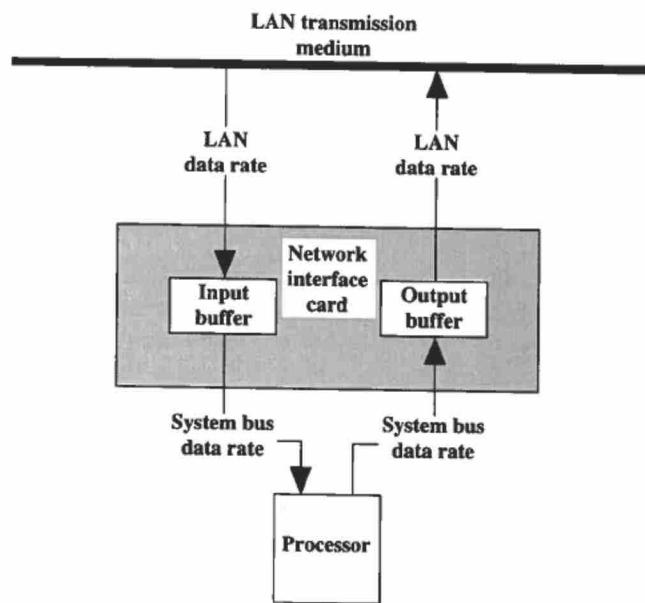


Figure 1.3 Buffered Transmission through a Network Interface Card

with the proliferation of devices, the growth in multimedia applications, and the increased use of the client/server architecture. As a result, much of the effort in LAN development has been in the development of high-speed LANs, with data rates of 100 Mbps or more. In later chapters, we will see a number of examples of high-speed LANs.

Metropolitan Area Networks

As the name suggests, a MAN occupies a middle ground between LANs and WANs. Interest in MANs has come about as a result of a recognition that the traditional point-to-point and switched network techniques used in WANs may be inadequate for the growing needs of organizations. While frame relay and ATM promise to meet a wide range of high-speed needs, there is a requirement now for both private and public networks that provide high capacity at low costs over a large area. The high-speed shared-medium approach of the LAN standards provides a number of benefits that can be realized on a metropolitan scale. As Figure 1.1 indicates, MANs cover greater distances at higher data rates than LANs, although there is some overlap in geographical coverage.

The primary market for MANs is the customer that has high-capacity needs in a metropolitan area. A MAN is intended to provide the required capacity at lower cost and greater efficiency than obtaining an equivalent service from the local telephone company.

1.3 APPLICATIONS OF LANs AND MANs

The variety of applications for LANs and MANs is wide. To provide some insight into the types of requirements that LANs and MANs are intended to meet, this section provides a brief discussion of some of the most important general application areas for these networks.

Personal Computer Local Networks

A common LAN configuration is one that supports personal computers. With the relatively low cost of such systems, individual managers within organizations often independently procure personal computers for departmental applications, such as spreadsheet and project management tools, and Internet access.

But a collection of department-level processors will not meet all of an organization's needs; central processing facilities are still required. Some programs, such as econometric forecasting models, are too big to run on a small computer. Corporate-wide data files, such as accounting and payroll, require a centralized facility but should be accessible to a number of users. In addition, there are other kinds of files that, although specialized, must be shared by a number of users. Further, there are sound reasons for connecting individual intelligent workstations not only to a central facility but to each other as well. Members of a project or organization team need to share work and information. By far the most efficient way to do so is digitally.

8 CHAPTER 1 / INTRODUCTION

Certain expensive resources, such as a disk or a laser printer, can be shared by all users of the departmental LAN. In addition, the network can tie into larger corporate network facilities. For example, the corporation may have a building-wide LAN and a wide area private network. A communications server can provide controlled access to these resources.

LANs for the support of personal computers and workstations have become nearly universal in organizations of all sizes. Even those sites that still depend heavily on the mainframe have transferred much of the processing load to networks of personal computers. Perhaps the prime example of the way in which personal computers are being used is to implement client/server applications.

For personal computer networks, a key requirement is low cost. In particular, the cost of attachment to the network must be significantly less than the cost of the attached device. Thus, for the ordinary personal computer, an attachment cost in the hundreds of dollars is desirable. For more expensive, high-performance workstations, higher attachment costs can be tolerated. In any case, this suggests that the data rate of the network may be limited; in general, the higher the data rate, the higher the cost.

Backend Networks and Storage Area Networks

Backend networks are used to interconnect large systems such as mainframes, supercomputers, and mass storage devices. The key requirement here is for bulk data transfer among a limited number of devices in a small area. High reliability is generally also a requirement. Typical characteristics include the following:

- **High data rate:** To satisfy the high-volume demand, data rates of 100 Mbps or more are required.
- **High-speed interface:** Data transfer operations between a large host system and a mass storage device are typically performed through high-speed parallel I/O interfaces, rather than slower communications interfaces. Thus, the physical link between station and network must be high speed.
- **Distributed access:** Some sort of distributed medium access control (MAC) technique is needed to enable a number of devices to share the medium with efficient and reliable access.
- **Limited distance:** Typically, a backend network will be employed in a computer room or a small number of contiguous rooms.
- **Limited number of devices:** The number of expensive mainframes and mass storage devices found in the computer room generally numbers in the tens of devices.

Typically, backend networks are found at sites of large companies or research installations with large data processing budgets. Because of the scale involved, a small difference in productivity can mean millions of dollars.

Consider a site that uses a dedicated mainframe computer. This implies a fairly large application or set of applications. As the load at the site grows, the existing mainframe may be replaced by a more powerful one, perhaps a multiprocessor system. At some sites, a single-system replacement will not be able to keep up;

Part Two surveys the key technology elements that are common to all types of LANs and MANs, including topology, transmission medium, medium access control, and logical link control.

CHAPTER 4 TOPOLOGIES AND TRANSMISSION MEDIA

The essential technology underlying all forms of LANs and MANs comprises topology, transmission medium, and medium access control technique. Chapter 4 examines the first two of these elements. Four topologies are in common use: bus, tree, ring, and star. The most common transmission media for local networking are twisted pair (unshielded and shielded), coaxial cable (baseband and broadband), optical fiber, and wireless. The chapter closes with a discussion of structured cabling systems.

CHAPTER 5 PROTOCOL ARCHITECTURE

Chapter 5 introduces the protocols needed for stations attached to a LAN to cooperate with each other in the exchange of data. Specifically, the chapter provides an overview of link control and medium access control protocols. The use of bridges and routers to interconnect LANs is also introduced.

CHAPTER 6 LOGICAL LINK CONTROL

Logical link control (LLC) is the highest layer that is specifically part of the LAN/MAN protocol architecture. It is used above all of the medium access control (MAC) standards. The primary purpose of this layer is to provide a means of exchanging data between end users across a link or a collection of LANs interconnected by bridges. Different forms of the LLC service are specified to meet specific reliability and efficiency needs. After a discussion of these services, Chapter 6 deals with some of the key mechanisms of link control protocols. Finally, the specific LLC protocols are examined.